

Portable Test and Stimulus Standard
Version 3.1

November 2026

Copyright © 2026 Accellera. All rights reserved.

Abstract: The definition of the language syntax and accompanying semantics for the specification of verification intent and behaviors reusable across multiple target platforms and allowing for the automation of test generation is provided. This standard provides a declarative environment designed for abstract behavioral description using actions, their inputs, outputs, and resource dependencies, and their composition into use cases including data and control flows. These use cases capture verification intent that can be analyzed to produce a wide range of possible legal scenarios for multiple execution platforms. It also includes a preliminary mechanism to capture the programmer's view of a peripheral device, independent of the underlying platform, further enhancing portability.

Keywords: behavioral model, constrained randomization, functional verification, hardware-software interface, portability, PSS, test generation.

Copyright © 2026 Accellera. All rights reserved.

Notices

Accellera Systems Initiative (Accellera) Standards documents are developed within Accellera and the Technical Committee of Accellera. Accellera develops its standards through a consensus development process, approved by its members and board of directors, which brings together volunteers representing varied viewpoints and interests to achieve the final product. Volunteers are members of Accellera and serve without compensation. While Accellera administers the process and establishes rules to promote fairness in the consensus development process, Accellera does not independently evaluate, test, or verify the accuracy of any of the information contained in its standards.

Use of an Accellera Standard is wholly voluntary. Accellera disclaims liability for any personal injury, property or other damage, of any nature whatsoever, whether special, indirect, consequential, or compensatory, directly or indirectly resulting from the publication, use of, or reliance upon this, or any other Accellera Standard document.

Accellera does not warrant or represent the accuracy or content of the material contained herein, and expressly disclaims any express or implied warranty, including any implied warranty of merchantability or suitability for a specific purpose, or that the use of the material contained herein is free from patent infringement. Accellera Standards documents are supplied **"AS IS."**

The existence of an Accellera Standard does not imply that there are no other ways to produce, test, measure, purchase, market, or provide other goods and services related to the scope of an Accellera Standard. Furthermore, the viewpoint expressed at the time a standard is approved and issued is subject to change due to developments in the state of the art and comments received from users of the standard. Every Accellera Standard is subjected to review periodically for revision and update. Users are cautioned to check to determine that they have the latest edition of any Accellera Standard.

In publishing and making this document available, Accellera is not suggesting or rendering professional or other services for, or on behalf of, any person or entity. Nor is Accellera undertaking to perform any duty owed by any other person or entity to another. Any person utilizing this, and any other Accellera Standards document, should rely upon the advice of a competent professional in determining the exercise of reasonable care in any given circumstances.

Interpretations: Occasionally questions may arise regarding the meaning of portions of standards as they relate to specific applications. When the need for interpretations is brought to the attention of Accellera, Accellera will initiate action to prepare appropriate responses. Since Accellera Standards represent a consensus of concerned interests, it is important to ensure that any interpretation has also received the concurrence of a balance of interests. For this reason, Accellera and the members of its Technical Committees are not able to provide an instant response to interpretation requests except in those cases where the matter has previously received formal consideration.

Comments for revision of Accellera Standards are welcome from any interested party, regardless of membership affiliation with Accellera. Suggestions for changes in documents should be in the form of a proposed change of text, together with appropriate supporting comments. Comments on standards and requests for interpretations should be addressed to:

Accellera Systems Initiative.
8698 Elk Grove Blvd Suite 1, #114
Elk Grove, CA 95624
USA

Note: Attention is called to the possibility that implementation of this standard may require use of subject matter covered by patent rights. By publication of this standard, no position is taken with respect to the existence or validity of any patent rights in connection therewith. Accellera shall not

Copyright © 2026 Accellera. All rights reserved.

be responsible for identifying patents for which a license may be required by an Accellera standard or for conducting inquiries into the legal validity or scope of those patents that are brought to its attention.

Accellera is the sole entity that may authorize the use of Accellera-owned certification marks and/or trademarks to indicate compliance with the materials set forth herein.

Authorization to photocopy portions of any individual standard for internal or personal use must be granted by Accellera, provided that permission is obtained from and any required fee is paid to Accellera. To arrange for authorization please contact Lynn Garibaldi, Accellera Systems Initiative, 8698 Elk Grove Blvd Suite 1, #114, Elk Grove, CA 95624, phone (916) 670-1056, e-mail lynn@accellera.org. Permission to photocopy portions of any individual standard for educational classroom use can also be obtained from Accellera.

Suggestions for improvements to the Portable Test and Stimulus Standard 3.1 are welcome. They should be posted to the PSS Community Forum at:

https://forums.accellera.org/forum/44-portable-stimulus-discussion/

The current Working Group web page is:

http://www.accellera.org/activities/working-groups/portable-stimulus

Copyright © 2026 Accellera. All rights reserved.

Introduction

The definition of a Portable Test and Stimulus Standard (PSS) will enable user companies to select the best tool(s) from competing vendors to meet their verification needs. Creation of a specification language for abstract use-cases is required. The goal is to allow stimulus and tests, including coverage and results checking, to be specified at a high level of abstraction, suitable for tools to interpret and create scenarios and generate implementations in a variety of languages and tool environments, with consistent behavior across multiple implementations.

This revision adds new features, corrects errors, clarifies aspects of the language and semantic definitions, removes some features, and reorganizes some sections relative to version 3.0 of the Portable Test and Stimulus Standard (August 2024).

The new features include (by section number):

Section(s)	Description
4.7.1	Added support for formatting string literals in a solve context
7.13	Added support for implementation- and model-level annotations elements
8.5.9 13.1.10	Added support for array-slice syntax with "in" and "unique" constraints
8.7	Clarified descriptions of numeric expressions treatment, with regard to type conversions and mixing different types
9.1.5	Added support for designating “instance” semantics for component ref fields
9.1.6	Added support for mutable attributes in components
9.2.2	Added support for action overriding in component subtypes
11.3.1.2	Added support for inline specification of non-rand action attribute values
11.3.2	Added support for multi-dimensional arrays of action handles
13.1.2	Added support for generic constraints
13.1.12	Added support for soft constraints
Clause 4 Clause 20	Clarified and enhanced target-template and triple-quotes
20.4.2	Added support for exported functions
20.5.4	Added exec block tags for code generation deduplication
21.1.3	Added support for “message” on both solve and target platforms
21.8	Added support for grouping realizations of different executor instances
21.9	Added support for target-time communication and synchronization between executors
21.11	Provided syntax to enable sharing an address claim between overlapping actions
21.13.4	Added support for negative offsets with “make_handle_from_claim” and “make_handle_from_handle”

Copyright © 2026 Accellera. All rights reserved.

Section(s)	Description
21.13.9	Added support for tools or users to extend the core library `mem_access_desc_s {}` to provide additional controls to the access functions. This additional memory access descriptor argument is not backward compatible.
21.14	Added register base types to support working with register collections
21.14.6	Added support for symbolic representation of registers in target code

Copyright © 2026 Accellera. All rights reserved.

Participants

The Portable Stimulus Working Group (PSWG) is entity-based. At the time this standard was developed, the PSWG had the following active participants:

Matthew Ballance, AMD, *Chair*
Tom Anderson, AMIQ EDA, *Secretary*
Jeanne Foster, *Technical Editor*
Shalom Bresticker, *Technical Editor Emeritus*

Agnisys, Inc.: Freddy Nunez

AMD: Matthew Ballance, Prabhat Gupta

AMIQ EDA: Tom Anderson

Arm Limited: Santosh Kumar, Karthikeyan Sugumaran

Big Fish EDA Consulting: Tom Fitzpatrick

Breker Verification Systems, Inc.: Leigh Brady

Cadence Design Systems, Inc.: Sergey Khaikin, Angelina Silver, Yuri Tsoglin

Ericsson AB: Ole Kristoffersen

Infineon Technologies AG: Leo Matturi

Intel Corporation: Jonathan Edwards

Microsoft Corporation: Scott Frazer

Qualcomm Technologies, Inc.: Lalana Palwaye, Rakshithraj Pulluri, Arjun Ashok Vazhayil

Synopsys, Inc.: Hillel Miller

At the time of standardization, the PSWG had the following eligible voters:

Agnisys, Inc.	**Cadence Design Systems, Inc.**
Advanced Micro Devices, Inc.	**Infineon Technologies AG**
AMIQ EDA	**Qualcomm Technologies, Inc.**
Arm Limited	**Synopsys, Inc.**

Copyright © 2026 Accellera. All rights reserved.

Contents

List of figures 19
List of tables 21
List of syntax excerpts 22
List of examples 26

1. Overview 34

1.1 Purpose 34
1.2 Language design considerations 34
1.3 Modeling basics 35
1.4 Test realization 35
1.5 Conventions used 36
1.5.1 Visual cues (meta-syntax) 36
1.5.2 BNF syntax conventions 37
1.5.3 Notational conventions 37
1.5.4 Examples 38
1.6 Use of color in this standard 38
1.7 Contents of this standard 38

2. References 39

3. Definitions, acronyms, and abbreviations 40

3.1 Definitions 40
3.2 Acronyms and abbreviations 41

4. Lexical conventions 42

4.1 Comments 42
4.2 Identifiers 42
4.3 Escaped identifiers 42
4.4 Keywords 43
4.5 Operators 43
4.6 Numbers 44
4.6.1 Integer constants 45
4.6.2 Floating-point constants 46
4.7 String literals 46
4.7.1 Special elements in triple-quoted string literals 48
4.7.2 Examples 51
4.8 Aggregate literals 52
4.8.1 Empty aggregate literal 53
4.8.2 Value list literals 53
4.8.3 Map literals 53
4.8.4 Structure literals 54
4.8.5 Nesting aggregate literals 55

5. Modeling concepts 56

5.1 Modeling data flow 57
5.1.1 Buffers 57
5.1.2 Streams 58
5.1.3 States 58

Copyright © 2026 Accellera. All rights reserved.

5.1.4 Data flow object pools ... 59
5.2 Modeling system resources ... 59
5.2.1 Resource objects ... 59
5.2.2 Resource pools ... 59
5.3 Basic building blocks ... 60
5.3.1 Components and binding ... 60
5.3.2 Evaluation and inference ... 60
5.4 Constraints and inferencing ... 62
5.5 Summary ... 62

6. Execution semantic concepts ... 63

6.1 Overview ... 63
6.2 Assumptions of abstract scheduling ... 63
6.2.1 Starting and ending action executions ... 63
6.2.2 Concurrency ... 63
6.2.3 Synchronized invocation ... 63
6.3 Scheduling concepts ... 64
6.3.1 Preliminary definitions ... 64
6.3.2 Sequential scheduling ... 65
6.3.3 Parallel scheduling ... 65
6.3.4 Concurrent scheduling ... 65
6.4 Test realization scheduling assumptions ... 65

7. Data types ... 66

7.1 General ... 66
7.1.1 Syntax ... 66
7.2 Integer types ... 67
7.2.1 Syntax ... 67
7.2.2 Examples ... 68
7.3 Floating-point types ... 68
7.3.1 Syntax ... 68
7.3.2 Cross-platform results ... 69
7.4 Booleans ... 69
7.5 Enumeration types ... 69
7.5.1 Syntax ... 69
7.5.2 Examples ... 71
7.6 Strings ... 72
7.6.1 Syntax ... 72
7.6.2 The sub-string operator ... 72
7.6.3 String methods ... 72
7.6.4 Example ... 73
7.7 Chandles ... 75
7.7.1 Syntax ... 76
7.7.2 Example ... 76
7.8 Structs ... 77
7.8.1 Syntax ... 77
7.8.2 Example ... 78
7.9 Collections ... 78
7.9.1 Syntax ... 78
7.9.2 Arrays ... 79
7.9.3 Lists ... 83
7.9.4 Maps ... 87

Copyright © 2026 Accellera. All rights reserved.

7.9.5 Sets ... 90
7.10 Reference types ... 93
7.10.1 Syntax ... 94
7.10.2 Examples ... 95
7.11 User-defined data types ... 97
7.11.1 Syntax ... 97
7.11.2 Example ... 97
7.12 Data type conversion ... 97
7.12.1 Syntax ... 97
7.12.2 Examples ... 98
7.13 Annotation types ... 99

8. Operators and expressions ... 102

8.1 Syntax ... 102
8.2 Constant expressions ... 102
8.3 Assignment operators ... 103
8.4 Expression operators ... 104
8.4.1 Operator precedence and associativity ... 104
8.4.2 Using aggregate literals in expressions ... 105
8.4.3 Type inference rules ... 107
8.4.4 Operator expression short-circuiting ... 108
8.5 Operator descriptions ... 109
8.5.1 Arithmetic operators ... 109
8.5.2 Relational operators ... 110
8.5.3 Equality operators ... 110
8.5.4 Logical operators ... 111
8.5.5 Bitwise operators ... 112
8.5.6 Reduction operators ... 113
8.5.7 Shift operators ... 113
8.5.8 Conditional operator ... 113
8.5.9 Set membership operator ... 114
8.6 Primary expressions ... 116
8.6.1 Bit-selects and part-selects ... 116
8.6.2 Selecting an element from a collection (indexing) ... 117
8.6.3 The sub-string operator ... 117
8.7 Evaluation rules for numeric expressions ... 117
8.7.1 Rules for expression types ... 118
8.7.2 Assignment-like contexts ... 120
8.7.3 Examples ... 122

9. Modeling element types ... 125

9.1 Components ... 125
9.1.1 Syntax ... 126
9.1.2 Example ... 126
9.1.3 Components as namespaces ... 127
9.1.4 Component instantiation ... 128
9.1.5 Component references ... 130
9.1.6 Mutable component attributes ... 134
9.1.7 Pure components ... 136
9.2 Actions ... 137
9.2.1 Syntax ... 138
9.2.2 Actions and inheritance ... 138

Copyright © 2026 Accellera. All rights reserved.

9.2.3 Examples ... 139
9.3 Flow objects ... 141
9.3.1 Buffer objects ... 141
9.3.2 Stream objects ... 142
9.3.3 State objects ... 143
9.3.4 Using flow objects ... 144
9.4 Resource objects ... 148
9.4.1 Declaring resource objects ... 148
9.4.2 Claiming resource objects ... 148

10. Template types ... 151

10.1 General ... 151
10.2 Template type declarations ... 151
10.2.1 Syntax ... 151
10.2.2 Example ... 152
10.3 Template parameter declarations ... 152
10.3.1 Template value parameter declarations ... 152
10.3.2 Template type parameter declarations ... 153
10.4 Template type instantiation ... 155
10.4.1 Syntax ... 155
10.4.2 Examples ... 156
10.5 Template type user restrictions ... 157

11. Action activities ... 158

11.1 Activity declarations ... 158
11.2 Activity constructs ... 158
11.2.1 Syntax ... 159
11.3 Action scheduling statements ... 159
11.3.1 Action traversal statement ... 159
11.3.2 Action handle array traversal ... 164
11.3.3 Sequential block ... 167
11.3.4 Parallel ... 168
11.3.5 Schedule ... 169
11.3.6 Fine-grained scheduling specifiers ... 172
11.3.7 Atomic block specifier ... 181
11.4 Activity control flow constructs ... 186
11.4.1 Repeat (count) ... 186
11.4.2 Repeat-while ... 187
11.4.3 Foreach ... 188
11.4.4 Select ... 189
11.4.5 If-else ... 191
11.4.6 Match ... 192
11.5 Activity construction statements ... 193
11.5.1 Replicate ... 193
11.6 Activity evaluation with extension and inheritance ... 197
11.7 Symbols ... 199
11.7.1 Syntax ... 199
11.7.2 Examples ... 200
11.8 Named sub-activities ... 201
11.8.1 Syntax ... 201
11.8.2 Scoping rules for named sub-activities ... 201
11.8.3 Hierarchical references using named sub-activity ... 202

Copyright © 2026 Accellera. All rights reserved.

11.9 Explicitly binding flow objects 204
11.9.1 Syntax 204
11.9.2 Examples 205
11.10 Hierarchical flow object binding 205
11.11 Hierarchical resource object binding 207

12. Pools 208

12.1 Syntax 208
12.2 Example 208
12.3 Static pool binding directive 209
12.3.1 Syntax 209
12.3.2 Examples 210
12.4 Resource pools and the instance_id attribute 214
12.5 Pool of states and the initial attribute 214

13. Randomization 216

13.1 Algebraic constraints 216
13.1.1 Member constraints 216
13.1.2 Generic constraints 218
13.1.3 Constraint inheritance 220
13.1.4 Action traversal in-line constraints 221
13.1.5 Logical expression constraints 222
13.1.6 Implication constraints 222
13.1.7 If-else constraints 223
13.1.8 Foreach constraints 225
13.1.9 Forall constraints 226
13.1.10 Unique constraints 229
13.1.11 Default value constraints 231
13.1.12 Soft constraints 237
13.1.13 Distribution directive 241
13.2 Scheduling constraints 243
13.2.1 Syntax 243
13.2.2 Example 244
13.3 Sequencing constraints on state objects 245
13.4 Randomization process 245
13.4.1 Random attribute fields 246
13.4.2 Randomization of lists 247
13.4.3 Randomization of flow objects 248
13.4.4 Randomization of resource objects 249
13.4.5 Randomization of component assignment 250
13.4.6 Procedural randomization of data 251
13.4.7 Random value selection order 253
13.4.8 Evaluation of expressions with action handles 253
13.4.9 Relationship lookahead 255
13.4.10 Lookahead and sub-actions 256
13.4.11 Lookahead and generic constraints 257
13.4.12 Pre_solve and post_solve exec blocks 258
13.4.13 Body blocks and sampling external data 261

14. Action inferencing 263

14.1 Implicit binding and action inferences 264

Copyright © 2026 Accellera. All rights reserved.

14.2 Object pools and action inferences 267
14.3 Data constraints and action inferences 269

15. Data coverage 271

15.1 Defining the coverage model: covergroup 271
15.1.1 Syntax 272
15.1.2 Examples 272
15.2 Covergroup instantiation 273
15.2.1 Syntax 274
15.2.2 Examples 274
15.3 Defining coverage points 275
15.3.1 Syntax 275
15.3.2 Example 276
15.3.3 Specifying bins 277
15.3.4 Automatic bin creation for coverage points 279
15.3.5 Excluding coverage point values 280
15.3.6 Specifying illegal coverage point values 280
15.3.7 Value resolution 281
15.4 Defining cross coverage 282
15.4.1 Syntax 282
15.4.2 Example 282
15.4.3 Defining cross bins 283
15.5 Specifying coverage options 284
15.5.1 Example 285
15.6 Covergroup sampling 285
15.7 Per-type and per-instance coverage collection 285
15.7.1 Per-instance coverage of flow and resource objects 286
15.7.2 Per-instance coverage in actions 286

16. Behavioral coverage 288

16.1 Defining behavioral coverage: cover and monitor 288
16.2 Behavioral coverage concepts 290
16.3 Monitor activity 293
16.3.1 Action traversal scenario 294
16.3.2 Sequential scenario 297
16.3.3 Concatenation scenario 298
16.3.4 Eventuality scenario 305
16.3.5 Overlapping scenario 307
16.3.6 Selection scenario 309
16.3.7 Empty scenario 311
16.3.8 Scheduling scenario 312
16.3.9 Monitor traversal 315
16.4 Monitor action handles and constraints 318
16.5 Covergroups in monitors 320
16.5.1 Covergroup sampling in monitors 320
16.5.2 Per-instance coverage in cover statements and monitors 324
16.6 Monitor activity evaluation with extension and inheritance 325

17. Type inheritance, extension, and overrides 327

17.1 Type inheritance 327
17.2 Type extension 334

Copyright © 2026 Accellera. All rights reserved.

17.2.1 Syntax ... 335
17.2.2 Examples ... 335
17.2.3 Struct and modeling element type extensions ... 336
17.2.4 Enumeration type extensions ... 337
17.2.5 Ordering of type extensions ... 338
17.2.6 Template type extensions ... 339
17.3 Combining inheritance and extension ... 340
17.4 Access protection ... 342
17.5 Overriding types ... 343
17.5.1 Syntax ... 344
17.5.2 Example ... 344

18. Source organization and processing ... 346

18.1 Packages ... 346
18.1.1 Package declarations ... 347
18.1.2 Nested packages ... 348
18.1.3 Referencing package members ... 349
18.1.4 Package aliases ... 351
18.2 Declaration and reference ordering ... 352
18.2.1 Examples ... 353
18.3 Name resolution ... 354
18.3.1 Name resolution examples ... 355

19. Conditional code processing ... 360

19.1 Overview ... 360
19.1.1 Statically-evaluated statements ... 360
19.1.2 Elaboration procedure ... 360
19.1.3 Compile-time expressions ... 360
19.2 Compile if ... 361
19.2.1 Scope ... 361
19.2.2 Syntax ... 362
19.2.3 Examples ... 362
19.3 Compile has ... 363
19.3.1 Syntax ... 363
19.3.2 Examples ... 364
19.4 Compile assert ... 365
19.4.1 Syntax ... 365
19.4.2 Example ... 365

20. Test realization ... 366

20.1 Exec blocks ... 366
20.1.1 Syntax ... 367
20.1.2 Exec block kinds ... 367
20.1.3 Examples ... 369
20.1.4 Exec block evaluation with inheritance and extension ... 372
20.1.5 Evaluation order of target exec blocks ... 375
20.2 Functions ... 379
20.2.1 Function declarations ... 380
20.2.2 Parameters and return types ... 382
20.2.3 Const parameters ... 383
20.2.4 Default parameter values ... 384

Copyright © 2026 Accellera. All rights reserved.

20.2.5 Generic and varargs parameters 385
20.2.6 Pure functions 386
20.2.7 Calling functions 386
20.3 Native PSS functions 390
20.3.1 Syntax 390
20.3.2 Parameter passing semantics 391
20.4 Foreign procedural interface 393
20.4.1 Definition using imported functions 393
20.4.2 Definition using exported functions 394
20.4.3 Imported classes 398
20.5 Target-template implementation of exec blocks 398
20.5.1 Target language 399
20.5.2 Exec file 399
20.5.3 Controlling the output of target-template exec blocks 399
20.5.4 Exec block tag 401
20.6 Target-template implementation for functions 405
20.6.1 Syntax 405
20.6.2 Examples 405
20.7 Procedural constructs 406
20.7.1 Scoped blocks 406
20.7.2 Variable declarations 407
20.7.3 Assignments 407
20.7.4 Void function calls 408
20.7.5 Return statement 409
20.7.6 Repeat (count) statement 409
20.7.7 Repeat-while statement 410
20.7.8 Foreach statement 411
20.7.9 If-else statement 412
20.7.10 Match statement 413
20.7.11 Break/continue statement 415
20.7.12 Randomize statement 416
20.7.13 Yield statement 417
20.7.14 Exec block 417
20.8 Blocking calls and concurrent execution 418
20.9 Comparison between mapping mechanisms 418
20.10 Exported actions 420
20.10.1 Syntax 420
20.10.2 Example 420
20.10.3 Export action foreign language binding 421

21. PSS core library 422

21.1 String formatting and output 422
21.1.1 String formatting 423
21.1.2 Solve-time string formatting and output 425
21.1.3 Message logging 425
21.2 File operations 426
21.3 Error reporting 429
21.4 Randomization 430
21.4.1 Urandom() 430
21.4.2 Urandom_range(min, max) 430
21.5 Floating-point 430
21.5.1 Floating-point storage types 431
21.5.2 Floating-point computation functions 431

Copyright © 2026 Accellera. All rights reserved.

21.5.3 Computation-type field extraction and composition ... 432
21.6 Standard annotations ... 433
21.6.1 Element code documentation ... 433
21.6.2 Element documentation ... 434
21.7 Executors ... 434
21.7.1 Executor representation ... 435
21.7.2 Executor assignment ... 437
21.8 Target execution units ... 445
21.9 Synchronization and communication ... 448
21.9.1 Channel type ... 449
21.10 Address spaces ... 450
21.10.1 Address space categories ... 450
21.10.2 Address space traits ... 453
21.10.3 Address space regions ... 455
21.11 Allocation within address spaces ... 456
21.11.1 Base claim type ... 456
21.11.2 Contiguous claims ... 456
21.11.3 Allocation mode ... 457
21.11.4 Transparent claims ... 460
21.11.5 Claim trait semantics ... 461
21.11.6 Allocation consistency ... 462
21.11.7 Rules for matching a claim to an address space ... 464
21.11.8 Allocation example ... 464
21.11.9 Address claim resolution ... 466
21.12 Address space group ... 466
21.12.1 Function add_addr_space ... 467
21.13 Data layout and access operations ... 470
21.13.1 Data layout ... 470
21.13.2 Sizeof_s ... 473
21.13.3 Address space handles ... 473
21.13.4 Obtaining an address space handle ... 474
21.13.5 Addr_value function ... 478
21.13.6 Addr_value_solve function ... 478
21.13.7 Addr_value_abs function ... 479
21.13.8 Get_tag function ... 479
21.13.9 Access operations ... 479
21.13.10 Target data structure setup example ... 490
21.14 Registers ... 492
21.14.1 PSS register definition ... 492
21.14.2 PSS register group definition ... 496
21.14.3 Association with address region ... 499
21.14.4 Obtaining register address handles ... 499
21.14.5 Translation of register read/write ... 500
21.14.6 Symbolic register names ... 501
21.14.7 Recommended packaging ... 508

Annex A
(informative)
Bibliography ... 509

Annex B
(normative)
Formal syntax ... 510

B.1 Package declarations ... 510

Copyright © 2026 Accellera. All rights reserved.

B.2 Action declarations ... 511
B.3 Struct declarations ... 513
B.4 Exec blocks ... 513
B.5 Functions ... 514
B.6 Foreign procedural interface ... 514
B.7 Procedural statements ... 515
B.8 Component declarations ... 516
B.9 Activity statements ... 517
B.10 Overrides ... 519
B.11 Data declarations ... 520
B.12 Template types ... 520
B.13 Data types ... 520
B.14 Constraints ... 522
B.15 Data coverage specification ... 523
B.16 Behavioral coverage specification ... 525
B.17 Conditional compilation ... 527
B.18 Expressions ... 528
B.19 Identifiers ... 529
B.20 Numbers and literals ... 530
B.21 Additional lexical conventions ... 532

Annex C
(normative)
Core library package ... 533

C.1 Package std_pkg ... 533
C.2 Package executor_pkg ... 535
C.3 Package addr_reg_pkg ... 535
C.4 Package sync_pkg ... 539

Annex D
(normative)
Foreign language bindings ... 541

D.1 Function prototype mapping ... 541
D.1.1 Export functions ... 541
D.2 Data type mapping ... 541
D.3 C language bindings ... 542
D.3.1 Function names ... 542
D.3.2 Primitive types ... 542
D.3.3 Arrays ... 542
D.3.4 Structs ... 543
D.3.5 Enumeration types ... 545
D.3.6 Export-function context handle ... 545
D.4 C++ language bindings ... 545
D.4.1 Function name mapping and namespaces ... 545
D.4.2 Primitive types ... 546
D.4.3 Arrays ... 546
D.4.4 Structs ... 547
D.4.5 Enumeration types ... 548
D.4.6 Export-function context handle ... 548
D.5 SystemVerilog language bindings ... 549
D.5.1 Function names ... 549
D.5.2 Primitive types ... 549
D.5.3 Numeric value mapping ... 550
D.5.4 Arrays ... 550

Copyright © 2026 Accellera. All rights reserved.

D.5.5 Lists ... 550
D.5.6 Structs ... 550
D.5.7 Enumeration types ... 551
D.5.8 Export-function context handle ... 551

Annex E
(informative)
Solution space ... 552

Annex F
(normative)
Formal semantics of behavioral coverage ... 554

F.1 General ... 554
F.2 Definitions and notation ... 554
F.3 Abstract syntax ... 555
F.3.1 Abstract grammar ... 555
F.3.2 Derived forms ... 555
F.4 Semantics ... 556
F.4.1 Action execution model ... 556
F.4.2 Scenario realization ... 557
F.4.3 Coverage semantics ... 558

Copyright © 2026 Accellera. All rights reserved.

List of figures

Figure 1—Partial specification of verification intent 56
Figure 2—Buffer flow object semantics 57
Figure 3—Stream flow object semantics 58
Figure 4—State flow object semantics 59
Figure 5—Single activity, multiple scenarios 61
Figure 6—Scheduling graph of activity with schedule block 171
Figure 7—Runtime behavior of activity with schedule block 171
Figure 8—Runtime behavior of scheduling block with sequential sub-blocks 172
Figure 9—join_branch scheduling graph 175
Figure 10—Join_branch runtime behavior 176
Figure 11—Scheduling graph of join_branch with scheduling dependency 177
Figure 12—Runtime behavior of join_branch with scheduling dependency 177
Figure 13—Join_none scheduling graph 178
Figure 14—Runtime behavior of join_first 179
Figure 15—Scheduling graph of join inside sequence block 180
Figure 16—Runtime behavior of join inside sequence block 180
Figure 17—Scheduling graph join with schedule block 181
Figure 18—Scheduling graph of action interference 183
Figure 19—Scheduling graph of atomic block avoiding interference 184
Figure 20—Scheduling graph of resource allocation issues 185
Figure 21—Monitor matching 292
Figure 22—Behavioral coverage concepts illustration 292
Figure 23—Action traversal statement. No match 296
Figure 24—Action traversal statement. One scenario realization 296
Figure 25—Matching action traversal statements with constraints 297
Figure 26—Action traversal, multiple matches 297
Figure 27—Sequential scenario 298
Figure 28—Action trace with two consecutive actions 299
Figure 29—Action trace with two repeated actions at the end 300
Figure 30—Action trace with two repeated actions in the middle 301
Figure 31—First alternative. Example 219 302
Figure 32—Second alternative. Example 219 302
Figure 33—First alternative. Example 220 303
Figure 34—Second alternative. Example 220 304
Figure 35—First alternative. Example 221 305
Figure 36—Second alternative. Example 221 305
Figure 37—First alternative. Example 222 306
Figure 38—Second alternative. Example 222 306
Figure 39—Overlapping scenario 308
Figure 40—Overlapping of three scenarios 309
Figure 41—Three scenarios. No overlap 309
Figure 42—Illustration to Example 225 311
Figure 43—Scheduling scenario 313
Figure 44—Scheduling scenario with common actions 314
Figure 45—Monitor traversal 316
Figure 46—Inlined and standalone constraints in monitors 320
Figure 47—Covergroup in a cover statement 321
Figure 48—Covergroup sampling. Multiple realizations 322
Figure 49—Covergroup instantiation in a monitor 323
Figure 50—Covergroup instantiation in a monitor. No match of cover statement 324
Figure 51—Order of invocation of init_down and init_up exec blocks 370

Copyright © 2026 Accellera. All rights reserved.

Figure 52—Address space regions with trait values 454
Figure 53—Different IP views of common storage atoms 468
Figure 54—Little-endian struct packing in register 472
Figure 55—Little-endian struct packing in byte-addressable space 472
Figure 56—Big-endian struct packing in register 472
Figure 57—Big-endian struct packing in byte-addressable space 473
Figure 58—Executor address mapping 488
Figure F.1—Action execution trace 557
Figure F.2—Action execution trace 559

Copyright © 2026 Accellera. All rights reserved.

List of tables

Table 1—Document conventions 36
Table 2—BNF syntax conventions 37
Table 3—PSS keywords 43
Table 4—Specifying special characters in string literals 47
Table 5—Control-flow directives in triple-quoted string literals 49
Table 6—Integer data types 67
Table 7—Floating-point computation data types 69
Table 8—Return type of sum() function 81
Table 9—Assignment operators and data types 103
Table 10—Expression operators and data types 104
Table 11—Operator precedence and associativity 104
Table 12—Binary arithmetic operators 109
Table 13—Power operator rules for integers 109
Table 14—Relational operators 110
Table 15—Equality operators 110
Table 16—Bitwise binary AND operator 112
Table 17—Bitwise binary OR operator 112
Table 18—Bitwise binary XOR operator 112
Table 19—Bitwise unary negation operator 112
Table 20—Results of unary reduction operations 113
Table 21—Types of primary expressions 118
Table 22—Types of compound expressions and their propagation to operands 119
Table 23—Assignment-like contexts 121
Table 24—Action handle array traversal contexts and semantics 166
Table 25—Instance-specific covergroup options 284
Table 26—Covergroup sampling 285
Table 27—Derived type element behaviors 328
Table 28—Flows supported for mapping mechanisms 419
Table 29—Exec block kinds supported for mapping mechanisms 419
Table 30—Data passing supported for mapping mechanisms 420
Table 31—Floating-point computation functions 431
Table 32—Scenario entity lifetimes 462
Table 33—Overlapping and sequential address claims examples 468
Table 34—Mnemonic construction for hierarchical register access 507
Table 35—Register access mode selection 507
Table D.1—Mapping PSS primitive types and C types 542
Table D.2—Mapping PSS struct types and C types 543
Table D.3—Mapping PSS struct field primitive types and C types 543
Table D.4—Mapping PSS enum types and C types 545
Table D.5—Mapping PSS primitive types and C++ types 546
Table D.6—Mapping PSS struct types and C++ types 547
Table D.7—Mapping PSS primitive types and SystemVerilog types 549

Copyright © 2026 Accellera. All rights reserved.

List of syntax excerpts

Syntax 1—Numeric constants 44
Syntax 2—Stringliterals 46
Syntax 3—Aggregate literals 52
Syntax 4—Empty aggregate literal 53
Syntax 5—Value list literal 53
Syntax 6—Map literal 53
Syntax 7—Structure literal 54
Syntax 8—Data types and data declarations 66
Syntax 9—Integer type declaration 67
Syntax 10—Floating-point type declaration 69
Syntax 11—Enum declaration 70
Syntax 12—String declaration 72
Syntax 13—Chandle declaration 76
Syntax 14—Struct declaration 77
Syntax 15—Collection data types 78
Syntax 16—Ref declaration 94
Syntax 17—User-defined type declaration 97
Syntax 18—Cast operation 97
Syntax 19—Annotation declaration syntax 100
Syntax 20—Annotation application syntax 100
Syntax 21—Expressions and operators 102
Syntax 22—Conditional operator 113
Syntax 23—Set membership operator 114
Syntax 24—String slice operator 117
Syntax 25—Component declaration 126
Syntax 26—Component ref instance qualifier 132
Syntax 27—Mutable qualifier 134
Syntax 28—Action declaration 138
Syntax 29—Buffer declaration 142
Syntax 30—Stream declaration 142
Syntax 31—State declaration 143
Syntax 32—Flow object reference 145
Syntax 33—Resource declaration 148
Syntax 34—Resource object reference 149
Syntax 35—Template type declaration 151
Syntax 36—Template value parameter declaration 152
Syntax 37—Template type parameter declaration 153
Syntax 38—Template type instantiation 155
Syntax 39—Activity statement 159
Syntax 40—Action traversal statement 160
Syntax 41—Activity sequence block 167
Syntax 42—Parallel statement 168
Syntax 43—Schedule statement 170
Syntax 44—Activity join specification 173
Syntax 45—Atomic block 182
Syntax 46—Repeat-count statement 186
Syntax 47—Repeat-while statement 187
Syntax 48—Foreach statement 188
Syntax 49—Select statement 189
Syntax 50—If-else statement 191
Syntax 51—Match statement 192

Copyright © 2026 Accellera. All rights reserved.

Syntax 52—Replicate statement 193
Syntax 53—Symbol declaration 199
Syntax 54—Bind statement 204
Syntax 55—Instantiation of pool 208
Syntax 56—Static bind directives 209
Syntax 57—Member constraint declaration 217
Syntax 58—Generic constraint declaration 218
Syntax 59—Expression constraint 222
Syntax 60—Implication constraint 222
Syntax 61—Conditional constraint 223
Syntax 62—Foreach constraint 225
Syntax 63—Forall constraint 226
Syntax 64—Unique constraint 229
Syntax 65—Default constraints 231
Syntax 66—Soft constraints 237
Syntax 67—Distribution directive 241
Syntax 68—Scheduling constraint statement 243
Syntax 69—Covergroup declaration 272
Syntax 70—Instantiating a covergroup 274
Syntax 71—Coverpoint declaration 276
Syntax 72—Bins declaration 277
Syntax 73—Cross declaration 282
Syntax 74—Cover statement and monitor declaration 288
Syntax 75—Monitor activity 293
Syntax 76—Action traversal statement 294
Syntax 77—Sequence monitor activity statement 297
Syntax 78—Concat monitor activity statement 298
Syntax 79—Eventually monitor activity statement 305
Syntax 80—Overlap monitor activity statement 307
Syntax 81—Select monitor activity statement 310
Syntax 82—Schedule monitor activity statement 312
Syntax 83—Monitor traversal statement 315
Syntax 84—Monitor constraints 318
Syntax 85—Type extension 335
Syntax 86—Override declaration 344
Syntax 87—Package declaration 347
Syntax 88—Import statement 350
Syntax 89—Compile if declaration 362
Syntax 90—Compile has expression 364
Syntax 91—Compile assert statement 365
Syntax 92—Exec block declaration 367
Syntax 93—Function declaration 380
Syntax 94—Function definition 390
Syntax 95—Imported function qualifiers 393
Syntax 96—Exported function 395
Syntax 97—Import class declaration 398
Syntax 98—Exec block tag 401
Syntax 99—Target-template function implementation 405
Syntax 100—Procedural block statement 406
Syntax 101—Procedural variable declaration 407
Syntax 102—Procedural assignment statement 408
Syntax 103—Void function call 409
Syntax 104—Procedural return statement 409
Syntax 105—Procedural repeat-count statement 410

Copyright © 2026 Accellera. All rights reserved.

Syntax 106—Procedural repeat-while statement ... 411
Syntax 107—Procedural foreach statement ... 412
Syntax 108—Procedural if-else statement ... 413
Syntax 109—Procedural match statement ... 414
Syntax 110—Procedural break/continue statement ... 415
Syntax 111—Procedural randomize statement ... 416
Syntax 112—Procedural yield statement ... 417
Syntax 113—Export action declaration ... 420
Syntax 114—String formatting and output functions ... 425
Syntax 115—Runtime messaging function ... 425
Syntax 116—Text file operations using file handles ... 427
Syntax 117—Simple text file operations ... 428
Syntax 118—Error reporting functions ... 429
Syntax 119—Randomization functions ... 430
Syntax 120—Floating-point storage types ... 431
Syntax 121—Float_mantissa function ... 432
Syntax 122—Float_exponent function ... 432
Syntax 123—Float_sign function ... 432
Syntax 124—To_float function ... 433
Syntax 125—Code_doc annotation ... 434
Syntax 126—Element documentation ... 434
Syntax 127—Executor component ... 436
Syntax 128—Executor group component ... 436
Syntax 129—Executor claim struct ... 438
Syntax 130—Executor query function ... 443
Syntax 131—Set_executor function ... 444
Syntax 132—Target_execution_unit_c ... 445
Syntax 133—PSS channel ... 449
Syntax 134—Generic address space component ... 450
Syntax 135—Contiguous address space component ... 451
Syntax 136—Transparent address space component ... 453
Syntax 137—Base address region type ... 455
Syntax 138—Contiguous address space region type ... 455
Syntax 139—Transparent region type ... 456
Syntax 140—Base address space claim type ... 456
Syntax 141—Contiguous address space claim type ... 457
Syntax 142—Allocation mode type ... 458
Syntax 143—Transparent contiguous address space claim type ... 461
Syntax 144—Address space group ... 467
Syntax 145—Packed_s base struct ... 471
Syntax 146—Sizeof_s struct ... 473
Syntax 147—Address space handle ... 474
Syntax 148—Make_handle_from_claim function ... 475
Syntax 149—Make_handle_from_handle function ... 476
Syntax 150—Addr_value function ... 478
Syntax 151—Addr_value_solve function ... 478
Syntax 152—Addr_value_abs function ... 479
Syntax 153—Get_tag function ... 479
Syntax 154—Mem_access_desc_s ... 480
Syntax 155—Primitive read operations for byte addressable spaces ... 483
Syntax 156—Primitive write operations for byte addressable spaces ... 483
Syntax 157—Read and write series of bytes ... 484
Syntax 158—Read and write packed structs ... 484
Syntax 159—Primitive operation implementation functions ... 485

Copyright © 2026 Accellera. All rights reserved.

Syntax 160—PSS register definition 492
Syntax 161—Reg_sized_c definition 493
Syntax 162—Reg_base_c definition 493
Syntax 163—PSS register group definition 497
Syntax 164—Mnemonic functions on reg_group_c 501
Syntax 165—Use_symbolic_reg_names function 502

Copyright © 2026 Accellera. All rights reserved.

List of examples

Example 1—Format-string expansion in static initialization 51
Example 2—Format-string expansion in procedural context 51
Example 3—Denoting multi- and single-line comments 52
Example 4—Value list literals 53
Example 5—Map literals 54
Example 6—Structure literals 54
Example 7—Nesting aggregate literals 55
Example 8—Enum data type 71
Example 9—String data type 73
Example 10—String operators and methods 75
Example 11—Chandle data type 76
Example 12—Struct with rand qualifiers 78
Example 13—Modifying collection contents 79
Example 14—Nested collection types 79
Example 15—Array declarations 80
Example 16—Fixed-size arrays 82
Example 17—Array operators and methods 83
Example 18—Declaring a list in a struct 83
Example 19—List operators and methods 86
Example 20—List randomization 87
Example 21—Declaring a map in a struct 87
Example 22—Map operators and methods 90
Example 23—Declaring a set in a struct 91
Example 24—Set operators and methods 93
Example 25—Use of reference as local variable and function return value 95
Example 26—Use of reference field and null value 96
Example 27—Typedef 97
Example 28—Overlap of possible enum values 98
Example 29—Casting of variable to an unsigned type 99
Example 30—Casting of reference type 99
Example 31—Annotation declaration 100
Example 32—Applying annotations 101
Example 33—Using a structure literal with an equality operator 106
Example 34—Using an aggregate literal with a set 106
Example 35—Using non-constant expressions in aggregate literals 106
Example 36—Contextual typing in structure literal interpretation 108
Example 37—Contextual typing in enum_item resolution 108
Example 38—Value range constraint 115
Example 39—Set membership in collection 115
Example 40—Set membership in variable range 115
Example 41—Set membership in array slice 116
Example 42—Bit size propagation 122
Example 43—Signedness propagation 122
Example 44—Expressions with integer and floating-point operands 123
Example 45—Type propagation with set membership operator 124
Example 46—Component 127
Example 47—Namespace 127
Example 48—Component declared in package 128
Example 49—Component instantiation 129
Example 50—Component attribute and function access 130
Example 51—Illegal traversal of an action outside of the containing component hierarchy 131

Copyright © 2026 Accellera. All rights reserved.

Example 52—Using the comp attribute in constraints 132
Example 53—Use of instance-qualified component ref fields 133
Example 54—Static bind and component instance-ref 134
Example 55—Mutable qualifier 136
Example 56—Pure components 137
Example 57—Override specification 139
Example 58—Atomic action 139
Example 59—Compound action 140
Example 60—Abstract action 140
Example 61—Overridden action 141
Example 62—Buffer object 142
Example 63—Stream object 143
Example 64—State object 144
Example 65—Buffer flow object 146
Example 66—Stream flow object 146
Example 67—Multiple producers/consumers using the same buffer pool 147
Example 68—Declaring a resource 148
Example 69—Resource object 150
Example 70—Locking and sharing arrays of resource objects 150
Example 71—Template type declarations 152
Example 72—Template value parameter declaration 153
Example 73—Another template value parameter declaration 153
Example 74—Template generic type and category type parameters 154
Example 75—Template parameter type restriction 154
Example 76—Template parameter used as base type 155
Example 77—Template type instantiation 156
Example 78—Template type qualification 156
Example 79—Overriding the default values 157
Example 80—Action traversal 162
Example 81—Anonymous action traversal 162
Example 82—Labeled action traversal 163
Example 83—Compound action traversal 163
Example 84—-Initialization at declaration and traversal levels 164
Example 85—Individual action handle array element traversal 165
Example 86—Action handle array traversal 165
Example 87—Two-dimensional action handle array traversal 166
Example 88—Sequential block 167
Example 89—Variants of specifying sequential execution in activity 168
Example 90—Parallel statement 169
Example 91—Another parallel statement 169
Example 92—Schedule statement 170
Example 93—Scheduling block with sequential sub-blocks 172
Example 94—Join_branch 174
Example 95—Join_branch with scheduling dependency 176
Example 96—Join_select 178
Example 97—Join_none 178
Example 98—Join_first 179
Example 99—Scope of join inside sequence block 179
Example 100—Specification for join with schedule block 181
Example 101—Atomic block to avoid action interference 183
Example 102—Atomic block to avoid resource allocation issues 185
Example 103—Repeat statement 186
Example 104—Another repeat statement 187
Example 105—Repeat-while statement 188

Copyright © 2026 Accellera. All rights reserved.

Example 106—Foreach statement ... 189
Example 107—Select statement ... 190
Example 108—Select statement with guard conditions and weights ... 191
Example 109—Select statement with array of action handles ... 191
Example 110—If-else statement ... 192
Example 111—Match statement ... 193
Example 112—Replicate statement ... 194
Example 113—Replicate statement with index variable ... 195
Example 114—Rewriting previous example without replicate statement ... 195
Example 115—Replicate statement with label array ... 196
Example 116—Replicate statement error situations ... 197
Example 117—Extended action traversal ... 198
Example 118—Hand-coded action traversal ... 198
Example 119—Inheritance and traversal ... 199
Example 120—Using a symbol ... 200
Example 121—Using a parameterized symbol ... 201
Example 122—Scoping and named sub-activities ... 202
Example 123—Label name conflict for activity statement ... 202
Example 124—Hierarchical references and named sub-activities ... 203
Example 125—Bind statement ... 205
Example 126—Hierarchical flow binding for buffer objects ... 206
Example 127—Hierarchical flow binding for stream objects ... 206
Example 128—Hierarchical resource binding ... 207
Example 129—Declaration of pool ... 208
Example 130—Static binding ... 210
Example 131—Binding of pools to array of components ... 211
Example 132—Pool binding ... 212
Example 133—Multiple state pools of the same state type ... 213
Example 134—Resource object assignment ... 214
Example 135—State object binding ... 215
Example 136—Declaring a static constraint ... 217
Example 137—Declaring a generic constraint ... 218
Example 138—Referencing a generic constraint inside a static constraint ... 219
Example 139—Value-yielding generic constraint ... 219
Example 140—Generic recursive constraint ... 220
Example 141—Inheriting and shadowing constraints ... 220
Example 142—Action traversal in-line constraint ... 221
Example 143—Name resolution inside **with** constraint block ... 222
Example 144—Implication constraint ... 223
Example 145—If constraint ... 224
Example 146—Foreach iterative constraint ... 226
Example 147—Forall constraint ... 227
Example 148—Rewrite of forall constraint in terms of explicit paths ... 228
Example 149—Forall constraint in different activity scopes ... 228
Example 150—Forall constraint item in a generic constraint ... 229
Example 151—Unique constraint ... 230
Example 152—Unique constraint on an array ... 230
Example 153—Unique constraint on a list slice ... 231
Example 154—Use of default value constraints ... 232
Example 155—Contradiction with default value constraints ... 233
Example 156—Default value constraints on compound data types ... 234
Example 157—Default constraints on flow objects ... 235
Example 158—Default constraints on flow objects with an explicit bind ... 236
Example 159—Default constraints on flow objects with an explicit hierarchical bind ... 237

Copyright © 2026 Accellera. All rights reserved.

Example 160—Soft-constraint priority 239
Example 161—Soft constraints and procedural randomization 240
Example 162—Aggregate soft constraints 240
Example 163—Distribution directive soft constraint 241
Example 164—Distribution directive on single variable 242
Example 165—Distribution directive on expression 242
Example 166—Distribution directive weight specification forms 243
Example 167—Constraint priority over distribution directive 243
Example 168—Zero-valued distribution weight 243
Example 169—Scheduling constraints 244
Example 170—Sequencing constraints 245
Example 171—Struct rand and non-rand fields 246
Example 172—Action rand-qualified fields 246
Example 173—Action-qualified fields 247
Example 174—Hierarchical constraint reference to list element 247
Example 175—Randomizing flow object attributes 249
Example 176—Randomizing resource object attributes 250
Example 177—Procedural randomization 251
Example 178—Evaluation of solve-time exec blocks in procedural randomization 252
Example 179—Activity with random fields 253
Example 180—Value selection of multiple traversals 254
Example 181—Illegal accesses to sub-action attributes 255
Example 182—Struct with random fields 255
Example 183—Activity with random fields 256
Example 184—Sub-activity traversal 257
Example 185—Activity with generic constraints 258
Example 186—Pre_solve/post_solve 260
Example 187—Post_solve ordering between action and flow objects 261
Example 188—Exec body block sampling external data 262
Example 189—Generating multiple scenarios 263
Example 190—Action inferences for partially-specified flows 265
Example 191—Buffer equality constraint to limit inferencing 266
Example 192—Resource equality constraint may affect scheduling 267
Example 193—Object pools affect inferencing 268
Example 194—Inferred traversal of an action outside of the containing component hierarchy 268
Example 195—In-line data constraints affect action inferencing 269
Example 196—Data constraints affect action inferencing 270
Example 197—Single coverage point 273
Example 198—Two coverage points and cross coverage items 273
Example 199—Creating and instantiating a covergroup type with a formal parameter list 274
Example 200—Creating a covergroup instance with instance-specific options 275
Example 201—Creating an in-line covergroup instance 275
Example 202—Specifying an iff condition 276
Example 203—Specifying bins 278
Example 204—Select constrained values between 0 and 255 279
Example 205—Using with in a coverpoint 279
Example 206—Excluding coverage point values 280
Example 207—Specifying illegal coverage point values 280
Example 208—Value resolution 281
Example 209—Specifying a cross 283
Example 210—Specifying cross bins 283
Example 211—Setting options 285
Example 212—Per-instance coverage of flow objects 286
Example 213—Per-instance coverage in actions 287

Copyright © 2026 Accellera. All rights reserved.

Example 214—Cover statement and monitor ... 290
Example 215—Illustration of behavioral coverage concepts ... 291
Example 216—Action traversal in monitors ... 295
Example 217—Concat vs sequence scenarios. Simple case ... 299
Example 218—Concat vs sequence scenarios. Covergroups ... 300
Example 219—Concat vs sequence scenarios. Inline constraints with same behavior ... 301
Example 220—Concat vs sequence scenarios. Different behavior due to standalone constraints ... 302
Example 221—Concat vs sequence scenarios. Inline constraints with different behavior ... 304
Example 222—Eventuality scenario ... 306
Example 223—Overlapping scenario ... 307
Example 224—Overlapping of three scenarios ... 309
Example 225—Illustration of behavioral coverage concepts ... 310
Example 226—Empty scenarios ... 311
Example 227—Degenerate scenario ... 312
Example 228—Scheduling scenario ... 313
Example 229—Scheduling scenario with common actions ... 314
Example 230—Monitor traversal ... 316
Example 231—Data constraint at monitor instantiation ... 317
Example 232—Action handles in monitors ... 319
Example 233—Inlined and standalone constraints in monitors ... 320
Example 234—Covergroup in a cover statement ... 321
Example 235—Covergroup sampling for multiple realizations ... 322
Example 236—Covergroup in a cover statement ... 323
Example 237—Per-instance coverage in monitors ... 325
Example 238—Monitor inheritance and traversal ... 326
Example 239—Declaring derived components and actions ... 329
Example 240—Default pool with inheritance ... 330
Example 241—Polymorphic function calls ... 331
Example 242—Derived type is also a base type ... 332
Example 243—Use of comp and this.comp with inheritance ... 333
Example 244—Illegal inheritance declaration ... 334
Example 245—Type extension ... 335
Example 246—Monitor type extension ... 336
Example 247—Action type extension ... 337
Example 248—Enum type extensions ... 338
Example 249—Template type extension ... 340
Example 250—Combining inheritance and extension ... 341
Example 251—Inheritance and extension of constraints ... 342
Example 252—Per-attribute access modifier ... 343
Example 253—Block access modifier ... 343
Example 254—Type inheritance and overrides ... 345
Example 255—Hierarchical declaration of nested package ... 348
Example 256—Direct declaration of nested package ... 348
Example 257—Declaration of nested package before outer package ... 348
Example 258—Importing the name of a nested package ... 351
Example 259—Package alias ... 352
Example 260—Illegal package alias declarations ... 352
Example 261—Reference to a previous source unit ... 353
Example 262—Reference to a later-declared action field ... 353
Example 263—Reference to local variable after declaration ... 353
Example 264—Initialization of constants ... 354
Example 265—Name resolution to declaration in nested namespace ... 355
Example 266—Name resolution to declaration in imported package in nested namespace ... 356
Example 267—Name resolution to declaration in encapsulating package ... 356

Copyright © 2026 Accellera. All rights reserved.

Example 268—Name resolution to declaration in imported package in encapsulating package 356
Example 269—Package import has no effect on name resolution 357
Example 270—Package import affects name resolution 357
Example 271—Package import is not a declaration 358
Example 272—Resolution of enum item references 358
Example 273—Resolution in presence of package alias 359
Example 274—Conditional compilation evaluation 361
Example 275—Conditional processing (C pre-processor) 363
Example 276—Conditional processing (compile if) 363
Example 277—Compile has 364
Example 278—Nested conditions 364
Example 279—Compile assert 365
Example 280—Data initialization in a component 369
Example 281—Init_down and init_up exec blocks 370
Example 282—Accessing component data field from an action 371
Example 283—Inheritance and shadowing 372
Example 284—Using super 373
Example 285—Type extension contributes an exec block 374
Example 286—Exec blocks added via extension 375
Example 287—Order of evaluation for run_start exec 377
Example 288—Function availability 382
Example 289—Reactive control flow 382
Example 290—Function declaration 383
Example 291—Const parameter declaration 384
Example 292—Default parameter value 385
Example 293—Generic parameter 385
Example 294—Varargs parameter 386
Example 295—Pure function 386
Example 296—Calling functions 387
Example 297—Function calls restrictions 389
Example 298—Parameter passing semantics 392
Example 299—Explicit specification of the implementation language 394
Example 300—Interrupt-driven PSS DMA model 396
Example 300—Interrupt-driven PSS DMA model (cont.) 397
Example 301—Environment interrupt-handling functions 397
Example 302—Import class 398
Example 303—Mustache expression used to specify a filename 399
Example 304—Referencing PSS variables using mustache notation 400
Example 305—Variable reference used to select the function 400
Example 306—Allowing programmatic declaration of target variable declarations 401
Example 307—Constant tag—generated once across instances 403
Example 308—Evaluated tag—generated per attribute value 403
Example 309—Tagged exec file in an executor component 403
Example 310—Exec tag with extension and inheritance 404
Example 311—Target-template function implementation 406
Example 312—Procedural return statement 409
Example 313—Procedural repeat-count statement 410
Example 314—Procedural while statement 411
Example 315—Procedural if-else statement 413
Example 316—Procedural match statement 414
Example 317—Procedural foreach statement with break/continue 416
Example 318—Exec block using procedural control flow statements 417
Example 319—Test-realization concurrency 418
Example 320—Export action 421

Copyright © 2026 Accellera. All rights reserved.

Example 321—Export action foreign language implementation 421
Example 322—Printing or formatting the context of a struct 425
Example 323—Runtime messages 426
Example 324—File operations 429
Example 325—Error reporting 430
Example 326—Conversion to and from storage type 433
Example 327—Applying the code_doc annotation in PSS 434
Example 328—Example of tool output from the code_doc annotation 434
Example 329—Defining an executor group 437
Example 330—Simple executor assignment 439
Example 331—Definition and use of executor trait 440
Example 332—Use of resource objects as executor claims 442
Example 333—Function delegation to executor 444
Example 334—Dual-cluster SOC example 447
Example 334—Dual-cluster SOC example (cont.) 448
Example 335—Contiguous address space in pss_top 452
Example 336—Example address trait type 453
Example 337—Address space with trait 454
Example 338—Shared claim with same address for embedded code using same compilation unit 459
Example 339—Overlapping shared claims 460
Example 340—Transparent address claim 461
Example 341—Address space allocation example 463
Example 342—Address space allocation example 465
Example 342—Address space allocation example (cont.) 466
Example 343—Address space group 469
Example 343—Address space group (cont.) 470
Example 344—Packed PSS little-endian struct 471
Example 345—Packed PSS big-endian struct 472
Example 346—Make_handle_from_claim example 475
Example 347—Negative offset usage 477
Example 347—Negative offset usage (cont.) 478
Example 348—Make_handle_from_handle example 478
Example 349—Customized access functions 481
Example 350—Customized access functions (frontdoor or backdoor) 482
Example 351—Illustration of read32() 486
Example 352—Mapping of primitive operations to foreign C functions 486
Example 353—Mapping of primitive operations to UVM sequences 487
Example 354—Implementing primitive operations in terms of other operations 487
Example 355—Customization of addr_value() 488
Example 356—Customization of memory functions 489
Example 357—Example using complex data structures 490
Example 357—Example using complex data structures (cont.) 491
Example 357—Example using complex data structures (cont.) 492
Example 358—Using sized register base 494
Example 359—Read-modify-write operations 495
Example 360—Examples of register declarations 496
Example 361—Example of register group declaration 498
Example 362—Top-level group and address region association 499
Example 363—Accessing registers based on their addresses 500
Example 364—Target platform symbolic register names 503
Example 365—Register group with flat symbolic names 504
Example 366—Target platform struct-based symbolic names 504
Example 367—Hierarchical register group with struct-accessor delimiter 505
Example 368—Full hierarchical example with symbolic name enable 506

Copyright © 2026 Accellera. All rights reserved.

Example 368—Full hierarchical example with symbolic name enable (Cont.) 507
Example 369—Recommended packaging 508
Example D.1—PSS struct mapping into C 544
Example D.2—PSS-defined exported function 545
Example D.3—C prototype for exported function 545
Example D.4—PSS-defined exported function 548
Example D.5—C prototype for exported function 549
Example D.6—PSS-defined exported function 551
Example D.7—SystemVerilog prototype for exported function 551

Copyright © 2026 Accellera. All rights reserved.

Portable Test and Stimulus Standard Version 3.1

1. Overview

This clause explains the purpose of this standard, describes its key concepts and considerations, details the conventions used, and summarizes its contents.

The Portable Test and Stimulus Standard syntax is specified using Backus-Naur Form (BNF). The rest of this standard is intended to be consistent with the BNF description. If any discrepancies between the two occur, the BNF formal syntax in Annex B shall take precedence.

1.1 Purpose

The Portable Test and Stimulus Standard defines a specification for creating a single representation of stimulus and test scenarios, usable by a variety of users across different levels of integration under different configurations, enabling the generation of different implementations of a scenario that run on a variety of execution platforms, including, but not necessarily limited to, simulation, emulation, FPGA prototyping, and post-silicon. With this standard, users can specify a set of behaviors once, from which multiple implementations may be derived.

1.2 Language design considerations

The Portable Test and Stimulus Standard (PSS) describes a declarative domain-specific language (DSL), intended for modeling scenario spaces of systems, generating test cases, and analyzing test runs. Scenario elements and formation rules are captured in a way that abstracts from implementation details and is thus reusable, portable, and adaptable. The portable stimulus specification captured in the DSL is herein referred to as *PSS*.

PSS borrows its core concepts from object-oriented programming languages, hardware-verification languages, and behavioral modeling languages. PSS features native constructs for system notions, such as data/control flow, concurrency and synchronization, resource requirements, and states and transitions. It also includes native constructs for mapping these to target implementation artifacts.

Introducing a new language has major benefits insofar as it expresses user intention that would be lost in other languages. However, user tasks that can be handled well enough in existing languages should be left to the language of choice, so as to leverage existing skill, tools, flows, and code bases. Thus, PSS focuses on

Copyright © 2026 Accellera. All rights reserved.

the essential domain-specific semantic layer and links with other languages to achieve other related purposes. This eases adoption and facilitates project efficiency and productivity.

Finally, PSS builds on prevailing linguistic intuitions in its constructs. In particular, its lexical and syntactic conventions come from the C/C++ family, and its constraint and coverage language uses SystemVerilog (IEEE Std 1800)[1] as a reference.

1.3 Modeling basics

A PSS *model* is a representation of some view of a system's behavior, along with a set of abstract flows. It is essentially a set of class definitions augmented with rules constraining their legal instantiation. A model consists of two types of class definitions: elements of behavior, called *actions*; and passive entities used by actions, such as resources, states, and data flow items, collectively called *objects*. The behaviors associated with an action are specified as *activities*. Actions and object definitions may be encapsulated in *components* to form reusable model pieces. All of these elements may also be encapsulated and extended in a *package* to allow for additional reuse and customization.

A particular instantiation of a given PSS model is called a *scenario*. Each scenario consists of a set of action instances and data object instances, as well as scheduling constraints and rules defining the relationships between them. The scheduling rules define a partial-order dependency relation over the included actions, which determines the execution semantics. A *consistent scenario* is one that conforms to model rules and satisfies all constraints.

Actions constitute the main abstraction mechanism in PSS. An action represents an element in the space of modeled behavior. Actions may correspond directly to operations of the underlying system under test (SUT) and test environment, in which case they are called *atomic actions*. Actions also use *activities* to encapsulate flows of simpler actions, constituting some joint activity or scenario intention. As such, actions can be used as top-level test intent or reusable test specification elements. Actions and objects have data attributes and data constraints over them.

Actions define the rules for legal combinations in general, not relative to a specific scenario. These are stated in terms of references to objects, having some role from the action's perspective. Objects thus serve as data, and control inputs and outputs of actions, or they are exclusively used as resources. Assembling actions and objects together, along with the scheduling and arithmetic constraints defined for them, produces a model that captures the full state-space of possible scenarios. A scenario is a particular solution of the constraints described by the model to produce an implementation consistent with the described intent.

1.4 Test realization

A key purpose of PSS is to automate the generation of test cases and test suites. Tests for electronic systems often involve code running on embedded controllers, exercising the underlying hardware and software layers. Tests may involve code in hardware-verification languages (HVLs) controlling bus functional models, as well as scripts, command files, data files, and other related artifacts. From the PSS model perspective, these are called *target files*, and *target languages*, which jointly implement the test case for a *target platform*.

The execution of a *consistent scenario* essentially consists of invoking its actions' implementations, if any, in their respective scheduling order. An action is invoked immediately after all its dependencies have completed, and subsequent actions wait for it to complete. Thus, actions that have the same set of

[1] Information on references can be found in Clause 2.

Copyright © 2026 Accellera. All rights reserved.

dependencies are logically invoked at the same time. Mapping atomic actions to their respective implementation for a target platform is captured in several ways, defined in Clause 20.

PSS features a native mechanism for referring to the actual state of the system under test (SUT) and the environment. Runtime values accessible to the generated test can be sampled and fed back into the model as part of an action's execution. These external values are sampled and, in turn, affect subsequent generation, which can be checked against model constraints and/or collected as coverage. The system/environment state can also be sampled during pre-run processing utilizing models and during post-run processing, given a run trace.

Similarly, the generation of a specific test-case from a given scenario may require further refinement or annotations, such as the external computation of expected results, memory modeling, and/or allocation policies. For these, external models, software libraries, or dedicated algorithmic code in other languages or tools may need to be employed. In PSS, the execution of these pre-run computations is defined using the same scheme as described above, with the results linked in the target language of choice.

1.5 Conventions used

The conventions used throughout the document are included here.

1.5.1 Visual cues (meta-syntax)

The meta-syntax for the description of the syntax rules uses the conventions shown in Table 1.

Table 1—Document conventions

Visual cue	Represents
bold	The **bold** font is used to indicate keywords and punctuation, text that shall be typed exactly as it appears. For example, in the following line, the keyword "state" and special characters "{" and "}" shall be typed exactly as they appear: **state** identifier [template_param_decl_list] [struct_super_spec] **{** { struct_body_item } **}**
plain text	The normal or plain text font indicates syntactic categories. For example, an identifier shall be specified in the following line (after the "state" keyword): **state** identifier [template_param_decl_list] [struct_super_spec] **{** { struct_body_item } **}**
italics	The *italics* font in running text indicates a definition. For example, the following line shows the definition of "activities": The behaviors associated with an action are specified as *activities*. The *italics* font in syntax definitions depicts a *meta-identifier*, e.g., *action*_identifier. See also 4.2.

Copyright © 2026 Accellera. All rights reserved.

Table 1—Document conventions (Continued)

Visual cue	Represents
`courier`	The `courier` font in running text indicates PSS code. For example, the following line indicates PSS code (for a `state`): `state power_state_s { int in [0..4] val; };`
{ } curly braces	Curly braces ({ }) indicate a set of action traversals. For example, the following sentence shows that "{start, write$_1$, read}" and "{start, write$_2$, read}" are action traversals. The top-level scenarios of `c5` and `c6` have the same realization for each trace: {start, write$_1$, read} for the trace in Figure 31 and {start, write$_2$, read} for the trace in Figure 32. See also 16.3.1.

1.5.2 BNF syntax conventions

The BNF syntax conventions are shown in Table 2.

Table 2—BNF syntax conventions

Visual cue	Represents
[ ] square brackets	Square brackets indicate optional items. For example, the *struct_super_spec* is optional in the following line: **state** identifier [template_param_decl_list] [struct_super_spec] **{** { struct_body_item } **}**
{ } curly braces	Curly braces ({ }) indicate items that can be repeated zero or more times. For example, the following line shows that zero or more *struct_body_item*s can be specified in this declaration: **state** identifier [template_param_decl_list] [struct_super_spec] **{** { struct_body_item } **}**
| separator bar	The separator bar (|) character indicates alternative choices. For example, the following line shows that the "input" or "output" keywords are possible values in a flow object reference: flow_ref_field_declaration ::= (**input** | **output**) flow_object_type object_ref_field { **,** object_ref_field } **;**
() parentheses	Parentheses (( )) group together alternative choices. For example, the following line shows that a flow object reference begins with either an "input" or an "output" keyword: flow_ref_field_declaration ::= (**input** | **output**) flow_object_type object_ref_field { **,** object_ref_field } **;**

1.5.3 Notational conventions

The terms "required", "shall", "shall not", "should", "should not", "recommended", "may", and "optional" in this document are to be interpreted as described in the IETF Best Practices Document 14, RFC 2119.

Copyright © 2026 Accellera. All rights reserved.

1.5.4 Examples

Any examples shown in this standard are for information only and are only intended to illustrate the use of PSS.

Many of the examples use "`. . .`" to indicate code omitted for brevity.

1.6 Use of color in this standard

This standard uses a minimal amount of color to enhance readability. The coloring is not essential and does not affect the accuracy of this standard when viewed in pure black and white. The places where color is used are the following:

- Cross references that are hyperlinked to other portions of this standard are shown in underlined-blue text (hyperlinking works when this standard is viewed interactively as a PDF file).
- Syntactic keywords and tokens in the formal language definitions are shown in **boldface-red text** when initially defined.

1.7 Contents of this standard

The organization of the remainder of this standard is as follows:

- Clause 2 provides references to other applicable standards that are assumed or required for this standard.
- Clause 3 defines terms and acronyms used throughout the different specifications contained in this standard.
- Clause 4 defines the lexical conventions used in PSS.
- Clause 5 defines the PSS modeling concepts.
- Clause 6 defines the PSS execution semantic concepts.
- Clause 7 highlights the PSS data types.
- Clause 8 describes the operators and operands that can be used in expressions and how expressions are evaluated.
- Clause 9 - Clause 18 describe the PSS abstract modeling constructs.
- Clause 19 describes the process for conditional code processing.
- Clause 20 describes the realization of PSS atomic actions.
- Clause 21 describes the PSS core library, which consists of portable functionality and utilities that PSS tools shall implement.
- Annexes. Following Clause 21 is a series of annexes.

Copyright © 2026 Accellera. All rights reserved.

2. References

The following referenced documents are indispensable for the application of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments or corrigenda) applies.

ANSI X3.4-1986: Coded Character Sets—7-Bit American National Standard Code for Information Interchange (7-Bit ASCII)[1] (ISO 646 International Reference Version)

IEEE Std 1800™-2023, IEEE Standard for SystemVerilog Unified Hardware Design, Specification and Verification Language.[2, 3]

The IETF Best Practices Document (for notational conventions) is available from the IETF web site: https://www.ietf.org/rfc/rfc2119.txt.

ISO/IEC 14882:2024, Programming Languages—C++ (commonly known as C++23).[4]

[1]ANSI publications are available from the American National Standards Institute (https://www.ansi.org/).

[2]The IEEE standards or products referred to in this clause are trademarks of the Institute of Electrical and Electronics Engineers, Inc.

[3]IEEE publications are available from the Institute of Electrical and Electronics Engineers, Inc., 445 Hoes Lane, Piscataway, NJ 08854, USA (https://standards.ieee.org/).

[4]ISO/IEC publications are available from the ISO Central Secretariat, Case Postale 56, 1 rue de Varembé, CH-1211, Genève 20, Switzerland/Suisse (https://www.iso.org/). ISO/IEC publications are also available in the United States from Global Engineering Documents, 15 Inverness Way East, Englewood, Colorado 80112, USA (https://global.ihs.com/). Electronic copies are available in the United States from the American National Standards Institute, 25 West 43rd Street, 4th Floor, New York, NY 10036, USA (https://www.ansi.org/).

Copyright © 2026 Accellera. All rights reserved.

3. Definitions, acronyms, and abbreviations

For the purposes of this document, the following terms and definitions apply. *The Authoritative Dictionary of IEEE Standards Terms* [B1][1] should be referenced for terms not defined in this clause.

3.1 Definitions

action: An element of behavior.

activity: An abstract, partial specification of a **scenario** that is used in a **compound action** or in a **compound monitor** to determine the high-level intent and leaves all other details open.

atomic action: An **action** that corresponds directly to operations of the underlying system under test (SUT) and test environment.

component: A structural entity, defined per type and instantiated under other components.

compound action: An **action** that includes an **activity** to traverse one or more sub-actions.

constraint: An algebraic expression relating attributes of model entities used to limit the resulting scenario space of the **model**.

coverage: A metric to measure the share of possible **scenario**s that have actually been processed for a given **model**.

exec block: Specifies the mapping of PSS scenario entities to their non-PSS implementation.

field: A variable associated with an instance of a type.

inheritance: The process of deriving one model element from another of a similar type, but adding or modifying functionality as desired. It allows multiple types to share functionality that only needs to be specified once, thereby maximizing reuse and portability.

loop: A traversal region of an **activity** in which a set of sub-actions is repeatedly executed. Values for the fields of the **action** are selected for each traversal of the loop, subject to the active constraints and resource requirements present.

model: A representation of some view of a system's behavior, along with a set of abstract flows.

monitor: An observed element of behavior.

object: A passive entity used by an **action**, such as resources, states, and data flow items.

override: To replace one or all instances of an element of a given type with an element of a compatible type inherited from the original type.

package: A way to group, encapsulate, and identify sets of related definitions, namely type declarations and type extensions.

resource: A computational element available in the target environment that may be claimed by an **action** for the duration of its execution.

[1]The numbers in brackets correspond to those of the bibliography in Annex A.

Copyright © 2026 Accellera. All rights reserved.

root action: An **action** designated explicitly as the entry point for the generation of a specific **scenario**. Any **action** in a **model** can serve as the root action of some **scenario**.

scenario: A particular instantiation of a given PSS model.

solve platform: The platform on which the test scenario is solved and, where applicable, target test code is generated. In some generation flows, the solve and target platforms may be the same.

target file: Contains textual content to be used in realizing the test intent.

target language: The language used to realize a specific unit of test intent, e.g., ANSI C, assembly language, Perl.

target platform: The execution platform on which test intent is executed.

type extension: The process of adding additional functionality to a model element of a given type, thereby maximizing reuse and portability. As opposed to **inheritance**, extension does not create a new type.

3.2 Acronyms and abbreviations

API Application Programming Interface

BNF Backus-Naur Form

ICL Inferencing Candidate List

MMIO Memory-Mapped Input/Output

PSS Portable Test and Stimulus Standard

SUT System Under Test

UVM Universal Verification Methodology

Copyright © 2026 Accellera. All rights reserved.

4. Lexical conventions

PSS borrows its lexical conventions from the C language family.

4.1 Comments

The token `/*` introduces a comment, which terminates with the first occurrence of the token `*/`. The C++ comment delimiter `//` is also supported and introduces a comment which terminates at the end of the current line.

4.2 Identifiers

An *identifier* is a sequence of letters, digits, and underscores, starting with a letter or an underscore; it is used to give an object a unique name so that it can be referenced. In a given namespace, identifiers shall be unique. Identifiers are case-sensitive.

A *meta-identifier* can appear in syntax definitions using the form: *construct_name*_identifier, e.g., *action*_identifier. See also B.19.

4.3 Escaped identifiers

Escaped identifiers shall start with the backslash character (`\`) and end with whitespace (space, tab, newline). They provide a means of including any of the printable non-whitespace ASCII characters in an identifier (the decimal values **`33`** through **`126`**, or **`0x21`** through **`0x7E`** in hexadecimal).

Neither the leading backslash character nor the terminating whitespace is considered to be part of the identifier. Therefore, an escaped identifier `\cpu3` is treated the same as a non-escaped identifier `cpu3`.

Some examples of legal escaped identifiers are shown here:

```
\busa+index
\-clock
\***error-condition***
\net1/\net2
\{a,b}
\a*(b+c)
```

Copyright © 2026 Accellera. All rights reserved.

4.4 Keywords

PSS reserves the keywords listed in Table 3.

Table 3—PSS keywords

abstract	**action**	**activity**	**annotation**	**array**	**as**
assert	**atomic**	**bind**	**bins**	**bit**	**body**
bool	**break**	**buffer**	**chandle**	**class**	**compile**
component	**concat**	**const**	**constraint**	**continue**	**cover**
covergroup	**coverpoint**	**cross**	**declaration**	**default**	**disable**
dist	**do**	**dynamic**	**else**	**enum**	**eventually**
exec	**export**	**extend**	**false**	**file**	**float32**
float64	**forall**	**foreach**	**function**	**has**	**header**
if	**iff**	**ignore_bins**	**illegal_bins**	**import**	**in**
init_down	**init_up**	**inout**	**input**	**instance**	**int**
join_branch	**join_first**	**join_none**	**join_select**	**list**	**lock**
map	**match**	**monitor**	**mutable**	**null**	**numeric**
output	**overlap**	**override**	**package**	**parallel**	**pool**
post_solve	**pre_body**	**pre_solve**	**private**	**protected**	**public**
pure	**rand**	**randomize**	**ref**	**repeat**	**replicate**
resource	**return**	**run_end**	**run_start**	**schedule**	**select**
sequence	**set**	**share**	**soft**	**solve**	**state**
static	**stream**	**string**	**struct**	**super**	**symbol**
target	**this**	**true**	**type**	**typedef**	**unique**
void	**while**	**with**	**yield**		

4.5 Operators

Operators are single-, double-, and triple-character sequences and are used in expressions. *Unary operators* appear to the left of their operand. *Binary operators* appear between their operands. A *conditional operator* has two operator characters that separate three operands.

Copyright © 2026 Accellera. All rights reserved.

4.6 Numbers

Constant numbers are specified as integer constants (see 4.6.1) or floating-point constants (see 4.6.2). The formal syntax for numbers is shown in Syntax 1.

```
number ::=
    integer_number
  | floating_point_number
integer_number ::=
    bin_number
  | oct_number
  | dec_number
  | hex_number
  | based_bin_number
  | based_oct_number
  | based_dec_number
  | based_hex_number
bin_digit ::= [0-1]
oct_digit ::= [0-7]
dec_digit ::= [0-9]
hex_digit ::= [0-9] | [a-f] | [A-F]
bin_number ::= 0[b|B] bin_digit { bin_digit | _ }
oct_number ::= 0 { oct_digit | _ }
dec_number ::= [1-9] { dec_digit | _ }
hex_number ::= 0[x|X] hex_digit { hex_digit | _ }
BASED_BIN_LITERAL ::= '[s|S]b|B bin_digit { bin_digit | _ }
BASED_OCT_LITERAL ::= '[s|S]o|O oct_digit { oct_digit | _ }
BASED_DEC_LITERAL ::= '[s|S]d|D dec_digit { dec_digit | _ }
BASED_HEX_LITERAL ::= '[s|S]h|H hex_digit { hex_digit | _ }
based_bin_number ::= [ dec_number ] BASED_BIN_LITERAL
based_oct_number ::= [ dec_number ] BASED_OCT_LITERAL
based_dec_number ::= [ dec_number ] BASED_DEC_LITERAL
based_hex_number ::= [ dec_number ] BASED_HEX_LITERAL
floating_point_number ::=
    floating_point_dec_number
  | floating_point_sci_number
unsigned_number ::= dec_digit { dec_digit | _ }
floating_point_dec_number ::= unsigned_number . unsigned_number
floating_point_sci_number ::=
    unsigned_number [ . unsigned_number ] exp [ sign ] unsigned_number
exp ::= e | E
sign ::= + | -
```

Syntax 1—Numeric constants

Copyright © 2026 Accellera. All rights reserved.

4.6.1 Integer constants

Integer literal constants can be specified in decimal, hexadecimal, octal, or binary format.

Several forms may be used to express an integer literal constant. The first form is a simple unsized decimal number, which is specified as a sequence of digits starting with **1** through **9** and containing the digits **0** through **9**.

The second form is an unsized hexadecimal number, which is specified with a prefix of **0x** or **0X** followed by a sequence of digits **0** through **9**, **a** through **f**, and **A** through **F**.

The third form is an unsized octal number, which is specified as a sequence of digits starting with **0** and containing the digits **0** through **7**.

The fourth form is an unsized binary number, which is specified with a prefix of **0b** or **0B** followed by a sequence of digits **0** and **1**.

The fifth form specifies a *based literal constant*, which is composed of up to three tokens:

- An optional size constant
- An apostrophe character (**'**) followed by a *base format* character
- Digits representing the value of the number.

The first token, a *size constant*, specifies the size of the integer literal constant in bits. This token shall be specified as an unsigned non-zero decimal number.

The second token, a *base format*, is a case-insensitive letter specifying the base for the number. The base is optionally preceded by the single character **s** (or **S**) to indicate a signed quantity. Legal base specifications are **d**, **D**, **h**, **H**, **o**, **O**, **b**, or **B**. These specify, respectively, decimal, hexadecimal, octal, and binary formats. The base format character and the optional sign character shall be preceded by an apostrophe. The apostrophe character and the base format character shall not be separated by whitespace.

The third token, an unsigned number, shall consist of digits that are legal for the specified base format. The unsigned number token immediately follows the base format, optionally separated by whitespace.

Simple decimal and octal numbers without the size and the base format shall be treated as *signed integers*. Unsized unbased hexadecimal and binary numbers shall be treated as unsigned. Numbers specified with a base format shall be treated as signed integers only if the **s** designator is included. If the **s** designator is not included, the number shall be treated as an unsigned integer.

If the size of an unsigned number is smaller than the size specified for the literal constant, the unsigned number shall be padded to the left with zeros. If the size of a signed number is smaller than the size specified for the literal constant, the signed number shall be sign-extended: padded to the left with zeros if the left-most bit of the original number is 0, and padded to the left with ones if the left-most bit of the original number is 1. If the size of a signed or unsigned number is larger than the size specified for the literal constant, the number shall be truncated from the left.

The number of bits that compose an unsized number shall be at least 32. An unsized number that requires more than 32 bits shall have the minimum width needed to properly represent the value, including a sign bit if the number is signed. For example, `0x7_0000_0000`, an *unsigned* hexadecimal number, shall have 35 bits. `4294967296` (2**32), a positive *signed* integer, shall be represented by 34 bits.

The underscore character (_) shall be legal anywhere in a number except as the first character. The underscore character can be used to break up long integer literals to improve readability.

Copyright © 2026 Accellera. All rights reserved.

When using integer literals in expressions, a negative value for an integer with no base specifier shall be interpreted differently from an integer with a base specifier. An integer with no base specifier shall be interpreted as a signed value in two's-complement form. An integer with an unsigned base specifier shall be interpreted as an unsigned value.

The following example shows four ways to write the expression "minus 12 divided by 3." Note that `-12` and `-'d12` both evaluate to the same two's-complement bit pattern, but, in an expression, the `-'d12` loses its identity as a signed negative number.

```
int IntA;
IntA = -12 / 3;         // The result is -4.
IntA = -'d12 / 3;       // The result is 1431655761.
IntA = -'sd12 / 3;      // The result is -4.
IntA = -4'sd12 / 3;     // -4'sd12 is the negative of the 4-bit quantity 1100,
                        // which is -4. -(-4) = 4. The result is 1.
```

4.6.2 Floating-point constants

Floating-point constant numbers can be specified either in decimal notation (e.g., `14.72`) or in scientific notation (e.g., `39e8`, which means 39 multiplied by 10 to the 8th power). Floating-point numbers expressed with a decimal point shall have at least one digit on each side of the decimal point. Whitespace is not permitted between the components of a floating-point constant.

Examples:

```
20.14     // Legal
20 .15    // Illegal. No whitespace is permitted between components.
2e6       // Legal, means 2 * 10**6
1e-9      // Legal, means 1 * 10**-9
```

4.7 String literals

A *string literal* is a sequence of ASCII characters enclosed by a single pair of quotation marks (`" ... "`), called a *quoted string*, or a triple pair of quotation marks (`""" ... """`), called a *triple-quoted string*. There is no predefined limit to the length of a string literal. The formal syntax for string literals is shown in Syntax 2.

```
string_literal ::=
    QUOTED_STRING
  | TRIPLE_QUOTED_STRING
QUOTED_STRING ::= " { unescaped_character | escaped_character } "
TRIPLE_QUOTED_STRING ::= """{any_ASCII_character}"""
unescaped_character ::= any_printable_ASCII_character
escaped_character ::= \('|"|?|\|a|b|f|n|r|t|v|[0-7][0-7][0-7])
```

Syntax 2—Stringliterals

PSS also includes a **string** data type to which a string literal can be assigned or compared. Variables of type **string** have arbitrary length; they are dynamically resized to hold any string. String literals are implicitly converted to the **string** type when assigned to a **string** type or used in an expression involving **string** type operands.

Copyright © 2026 Accellera. All rights reserved.

The *empty string literal* (`""`) represents an empty, or null, string.

Quoted string literals may only contain printable ASCII characters (the decimal values **32** through **126**, or **0x20** through **0x7E** in hexadecimal). Certain characters can be used in quoted string literals when preceded by an *escape character* (a *backslash*). Table 4 lists these characters, with the escape sequence that represents them. A quoted string shall be contained in a single line.

Table 4—Specifying special characters in string literals

Escape sequence	ASCII hex value	Character produced by escape sequence
\a	**0x07**	Alert (Beep, Bell)
\b	**0x08**	Backspace
\f	**0x0C**	Formfeed
\n	**0x0A**	Newline
\r	**0x0D**	Carriage return
\t	**0x09**	Horizontal tab
\v	**0x0B**	Vertical tab
****	**0x5C**	\ character (backslash)
\"	**0x22**	" character (double quotation mark)
\'	**0x27**	' character (apostrophe, single quotation mark)
\?	**0x3F**	? character (question mark)
\ddd	**any**	A character specified in 3 octal digits (see Syntax 1). Implementations may issue an error if the character represented is greater than **\377**.

An escape sequence is considered a single character in the string literal. An escaped apostrophe or question mark is treated the same as an unescaped apostrophe or question mark, respectively, i.e., the backslash is ignored. The other escaped characters in the table have different meanings from their unescaped versions. It is illegal for an escape character in a quoted string literal to be followed by any character not appearing in the table above.

In contrast, a triple-quoted string literal may contain any ASCII character, printing or nonprinting. There is no escape character. For example, triple-quoted strings may contain both single and double quotation marks (except for three consecutive double quotation marks) and newline characters.

In addition, a triple-quoted string literal may contain special elements, as described in 4.7.1. Such elements are processed according to the rules in 4.7.1. Characters that are not part of a special element are included in the resulting string as written, except where those rules explicitly specify otherwise.

Both quoted string literals and triple-quoted string literals may be used anywhere a string literal is desired or required, with the following exception: if a triple-quoted string contains special elements that depend on any non-constant expressions (e.g., non-constant mustache expressions or control-flow directives depending on non-constant expressions), it cannot be used where a constant string is expected.

Copyright © 2026 Accellera. All rights reserved.

4.7.1 Special elements in triple-quoted string literals

A triple-quoted string literal may contain the following special elements: mustache expressions, control-flow directives, and embedded comments. These can be used in:

— A target-template block (in the context of a target exec or a target function) or the filename in the context of an exec-file block. In this case, variables from the enclosing type may be referenced. Solve function calls may be used, provided they have no side effects; calling a function that has side effects from within a target-template block may cause unspecified results. Mustache expressions and control-flow directives are expanded after the PSS **pre-body** phase has completed.

— A string expression. In this case, whatever can be referenced in the same context not as part of a string literal, can also be referenced within a mustache expression or a control-flow directive. Mustache expressions and control-flow directives are expanded as part of the string expression evaluation.

A triple-quoted string that references only constant expressions via mustache notation or control-flow directives is, itself, a constant.

4.7.1.1 Mustache notation

Triple-quoted strings may contain expressions inside mustache notation (**`{{expression}}`**) that are expanded when the string is evaluated in specific contexts.

A mustache expression may reference variables visible from the scope in which the triple-quoted string is used. The expression type shall be a scalar type.

When the expression value is converted to a string, the result is formatted as follows:

— **int** signed decimal (**%d**)

— **bit** unsigned decimal (**%u**)

— **bool** "**true**" | "**false**" (**%n**)

— **enum** the string identifier associated with the enumerator for the value

— **string** string (**%s**)

— **chandle** pointer (**%p**)

— **float32**, **float64** floating-point (**%f**)

4.7.1.2 Control-flow directives

Triple-quoted strings may contain control-flow directives enclosed in **{% … %}** delimiters. The supported directives are listed in Table 5. Any expression used in these directives (for example, the collection expression iterated by a **foreach** loop directive) shall respect the following rules: it may reference variables from the scope in which the triple-quoted string is used. However, when used within control-flow directives, expressions shall not be enclosed in mustache notation.

No characters are added to the string from the control-flow directives themselves. These directives only affect the logic of adding following text into the string.

If a line within the triple-quoted string literal contains one or more control-flow directives, and no other characters other than whitespace are present on the same line, then these whitespace characters (including the new-line character ending the line) shall be excluded from the resulting string.

Copyright © 2026 Accellera. All rights reserved.

Table 5—Control-flow directives in triple-quoted string literals

Directive	Syntax	Description
if	**{% if (** expression **) %}**	This shall start a block, closed with a later **else**, **else if**, or ending **{%%}** directive. *expression* is a Boolean expression. It is evaluated, and if the condition holds (i.e., *expression* evaluates to **true**), the text within the block shall be added to the string, until the closing directive.
else	**{% else %}**	This shall end a block started by a previous **if** or **else if** directive, and shall start a new block, closed with a later ending **{%%}** directive. If the condition of the preceding **if** or **else if** directive was evaluated and did not hold, the text within the new block shall be added to the string, until the closing directive.
else if	**{% else if (** expression **) %}**	This shall end a block started by a previous **if** or **else if** directive. This shall start a new block, closed with a later **else**, **else if**, or ending **{%%}** directive. *expression* is a Boolean expression. If the conditions of the preceding **if** or **else if** directive were evaluated and did not hold, and the condition of this directive holds (i.e., *expression* evaluates to **true**), the text within the new block shall be added to the string, until the closing directive.
repeat	**{% repeat (**[*index*_identifier **:**] expression **) %}**	This directive shall start a block closed with a later ending **{%%}** directive. *expression* shall be a non-negative numeric expression. The text within the block shall be added repetitively, the number of times specified by *expression*. The index variable, if used, shall be of type **int** and range between 0 and one less than the number of iterations. The index variable is added to the scope until the matching ending directive, and can be referenced in mustache expressions and other control-flow directives within this scope; however, its value cannot be explicitly assigned to or otherwise changed within this scope.

Copyright © 2026 Accellera. All rights reserved.

Table 5—Control-flow directives in triple-quoted string literals (Continued)

Directive	Syntax	Description
foreach	**{% foreach (** [*iterator*_identifier **:**] expression [**[** *index*_identifier **]**] **) %}**	This shall start a block closed with a later ending **{%%}** directive. *expression* is a collection type expression. The text within the block shall be added repetitively, as per the number of elements in the collection *expression*. *iterator_identifier*, if used, specifies the name of an iterator variable of the collection element type. *index_identifier*, if used, specifies the name of an index variable, as follows: — If the collection is **array** or **list**, the index variable shall be of type **int**, ranging from **0** to one less than the size of the collection, in that order. — If the collection is **map**, the index variable shall be of the same type as the **map** keys, and range over the values of the keys. — If the collection is **set**, the index variable shall not be used. Either an index variable or an iterator variable or both shall be used. They shall be added to the scope until the block closing directive, and can be referenced in mustache expressions and other control-flow directives within this scope; however, their value cannot be explicitly assigned to or otherwise changed within this scope.
Ending	**{%%}**	This shall follow and match a previous **if**, **else**, **else if**, **foreach**, or **repeat** directive, ending the effect of the former.
Variable declaration	**{%** data_type procedural_data_instantiation { **,** procedural_data_instantiation } **; %}**	This shall introduce a local variable visible in the current scope, similar to declaring variables in procedural contexts (see 20.7.2). This variable can be referenced in mustache expressions and other control-flow directives within this scope.
Variable assignment	**{%** identifier **=** expression **; %}**	This shall assign the value of *expression* into the variable specified by *identifier*. This can only be used with variables previously declared with a variable declaration directive within the same triple-quoted string and visible in the current scope. This cannot be used with any other variables, such as loop variables introduced within the same triple-quoted string, or action attributes.

Copyright © 2026 Accellera. All rights reserved.

4.7.1.3 Embedded comments

To retain parts of a triple-quoted string literal as PSS comments and prevent the commented text from appearing in the resulting string (such as a target code, when used in target-template blocks) the language allows a special commenting notation with a hash inside braces (without whitespace). The token `{#` introduces a comment, which ends with the successive occurrence of `#}`, enabling a multi-line comments capture. And the notation `{#}` introduces a comment that ends with the current line.

Text within an embedded comment is not processed for special elements. In particular, any mustache expressions or control-flow directives appearing within an embedded comment shall not be expanded or interpreted.

4.7.2 Examples

The following string literals are equivalent:

```
" \"Humpty Dumpty sat on a wall.\nHumpty Dumpty had a great fall.\" "

""" "Humpty Dumpty sat on a wall.
Humpty Dumpty had a great fall." """
```

Example 1 demonstrates initialization of static constants with a string literal containing mustache expressions. The mustache expressions are evaluated when the static constant is instantiated and assigned a value. Therefore, `C2` is initialized with the string value "2".

```
static const int C1 = 2;
static const string C2 = """C1={{C1}}""";
static const string C3 = """C1={{C4}}"""; // illegal reference, since
                                          // C4 has not yet been
                                          // declared
```

Example 1—Format-string expansion in static initialization

In Example 2, mustache expressions in the triple-quoted string assigned to `v` are evaluated when the assignment is performed. At this point in time, `a=5`, so `v` is assigned the string "`55555`": the **repeat** directive having index `5`, each time adding "`5`" to the string. When function `f1` is called, `a=6`, so the string "`a=6`" is passed to the function `f1`.

```
function void f1(string val);

function void f2() {
  int a = 5;
  string v = """{% repeat (a) %}{{a}}{%%}"""; // v = "55555"
  a = 6;
  f1("""a={{a}}"""); // "a=6" is passed to f1
}
```

Example 2—Format-string expansion in procedural context

Example 3 demonstrates the usage of a multi-line comment where an implementation with an individual mode-based API declaration is commented and replaced with a declaration based on the action usage. The example also captures the usage of a single-line comment where a constant variable implementation is

Copyright © 2026 Accellera. All rights reserved.

commented and replaced with a static one. The mustache expression appearing in it is part of the comment text and is not expanded.

```
enum op_mode_e {rx,tx,copy};
component transactor {
  list <op_mode_e> op_mode_l;
  exec init_up {
    op_mode_l = {rx, tx};
  }

  action read_a {
    rand op_mode_e opmode;
    rand int trans_size;

    constraint opmode in comp.op_mode_l;

    exec declaration C = """
      //This comment will appear in the target code.
      //{{opmode}} - Solved value will appear in the target code.
      {#
         This comment block will not appear in target code.
         static void transactor_init_rx_init();
         static void transactor_init_tx_init();
         static void transactor_init_copy_init();
      #}

      static void transactor_init_{{opmode}}_init();

      {#} const int trans_size = {{trans_size}};
      static int trans_size = {{trans_size}};
    """;
  }
}
```

Example 3—Denoting multi- and single-line comments

4.8 Aggregate literals

Aggregate literals are used to specify the content values of collections and structure types. The different types of aggregate literals are described in the following sections. The use of aggregate literals in expressions is described in 8.4.2.

```
aggregate_literal ::=
    empty_aggregate_literal
  | value_list_literal
  | map_literal
  | struct_literal
```

Syntax 3—Aggregate literals

Copyright © 2026 Accellera. All rights reserved.

4.8.1 Empty aggregate literal

```
empty_aggregate_literal ::= { }
```

Syntax 4—Empty aggregate literal

Aggregate literals with no values specify an empty *collection* (see 7.9) when used in the context of a variable-sized collection type (**list**, **set**, **map**).

4.8.2 Value list literals

```
value_list_literal ::= { expression { , expression } }
```

Syntax 5—Value list literal

Aggregate literals for use with **array**s, **list**s, and **set**s (see 7.9) use *value list literal*s. Each element in the list specifies an individual value. When used in the context of **list**s, the number of elements in the value list literal specifies the size as well as the values. When used in the context of **set**s, the number of unique elements in the value list literal specifies the size as well as the values (equivalent elements are counted only once). When used in the context of **array**s and **list**s, the value list literal also specifies the order of elements, starting with element 0. The data types of the values shall match the data type specified in the collection declaration.

When a value list literal is used in the context of an **array**, the value list literal shall have the same number of elements as the **array**. It is an error if the value list literal has more or fewer elements than the **array**.

```
int c1[4] = {1, 2, 3, 4};       // OK
int c2[4] =  {1};               // Error: literal has fewer elements than array
int c3[4] = {1, 2, 3, 4, 5, 6}; // Error: literal has more elements than array
```

Example 4—Value list literals

Values in value list literals may be non-constant expressions.

If the value list literal elements are numbers, the exact element type of the collection is determined by context according to the expression type propagation rules specified in 8.7.1. For example, if used as the right-hand side of the assignment operator (see 8.3), the list elements type is the same as the left-hand side element type. Each value list literal element expression shall be implicitly cast to that type. If there is no context type (for example, the value list literal is used as the *expression* of a **foreach** statement), all elements shall be of the same self-determined type; otherwise, an error shall be generated. When used as the right-hand side of a set membership operator, a different type propagation rule is used (see Table 22).

4.8.3 Map literals

```
map_literal ::= { map_literal_item { , map_literal_item } }
map_literal_item ::= expression : expression
```

Syntax 6—Map literal

Copyright © 2026 Accellera. All rights reserved.

Aggregate literals for use with **map**s (see 7.9.4) use *map literals*. The first element in each colon-separated pair is the key. The second element is the value to be associated with the key. The data types of the expressions shall match the data types specified in the **map** declaration. If the same key appears more than once, the last value specified is used.

In Example 5, a map literal is used to set the value of a **map** with integer keys and Boolean values.

```
struct t {
  map<int,bool> m = {1:true, 2:false, 4:true, 8:false};
  constraint m[1]; // True, since the value "true" is associated with key "1"
}
```

Example 5—Map literals

Both keys and values in map literals may be non-constant expressions.

If the keys and/or values in a map literal are numbers, the exact key and/or value type of the map is determined by context according to the expression type propagation rules specified in 8.7.1. For example, if used as the right-hand side of the assignment operator (see 8.3), the map key and value types are the same as the corresponding left-hand side map key and value types, accordingly. Each key expression in the map literal shall be implicitly cast to the context map key type, and each value expression in the map literal shall be implicitly cast to the context map value type. If there is no context type (for example, the map literal is used as the *expression* of a **foreach** statement), all key elements shall be of the same self-determined type, and all value elements shall be of the same self-determined type; otherwise, an error shall be generated.

4.8.4 Structure literals

```
struct_literal ::= { struct_literal_item { , struct_literal_item } }
struct_literal_item ::= . identifier = expression
```

Syntax 7—Structure literal

A *structure literal* explicitly specifies the name of the **struct** attribute that a given expression is associated with. **Struct** attributes whose value is not specified are assigned the default value of the attribute. The order of the attributes in the literal does not have to match their order in the **struct** declaration. It shall be illegal to specify the same attribute more than once in the literal.

In Example 6, the initial value for the attributes of `s1` is explicitly specified for all attributes. The initial value for the attributes of `s2` is specified for a subset of attributes. The resulting value of both `s1` and `s2` is `{.a=1,.b=2,.c=0,.d=0}`. Consequently, the constraint `s1==s2` holds.

```
struct s {
  int a, b, c, d;
};
struct t {
  s s1 = {.a=1,.b=2,.c=0,.d=0};
  s s2 = {.b=2,.a=1};
  constraint s1 == s2;
}
```

Example 6—Structure literals

Copyright © 2026 Accellera. All rights reserved.

Values in structure literals may be non-constant expressions.

4.8.5 Nesting aggregate literals

Aggregate literals may be nested to form the value of data structures formed from nesting of aggregate data types.

In Example 7, an aggregate literal is used to form a list of **struct** values. Each structure literal specifies a subset of the **struct** attributes.

```
struct s {
  int a, b, c, d;
};
struct t {
  list<s>  my_l = {
    {.a=1,.d=4},
    {.b=2,.c=8}
  };
}
```

Example 7—Nesting aggregate literals

Copyright © 2026 Accellera. All rights reserved.

5. Modeling concepts

A PSS model is made up of a number of elements (described briefly in 1.3) that define a set of possible scenarios to be applied to the Design Under Test (DUT) via the associated test environment. *Scenarios* are composed of behaviors—ultimately executed on some combination of components that make up the DUT or on verification components that define the test environment—and the communication between them. This clause introduces the elements of a PSS model and defines their relationships.

The primary behavior abstraction mechanism in PSS is an *action*, which represents a particular behavior or set of behaviors. Actions combine to form the scenarios that represent the verification intent. Actions that correspond directly to operations performed by the underlying DUT or test environment are referred to as *atomic actions*, which contain an explicit mapping of the behavior to an implementation on the target platform in one of several supported forms. *Compound actions* encapsulate flows of other actions using an *activity* that defines the critical intent to be verified by specifying the relationships between specific actions.

The remainder of the PSS model describes a set of rules that are used by a PSS processing tool to create the scenarios that implement the critical verification intent while satisfying the data flow, scheduling, and resource constraints of the target DUT and associated test environment. In the case where the specification of intent is incomplete (partial), the PSS processing tool shall infer the execution of additional actions and other model elements necessary to make the partial specification complete and valid. In this way, a single partial specification of verification intent may be expanded into a variety of actual scenarios that all implement the critical intent, but might also include a wide range of other behaviors that may provide greater coverage of the functionality of the DUT as demonstrated in the example in Figure 1.

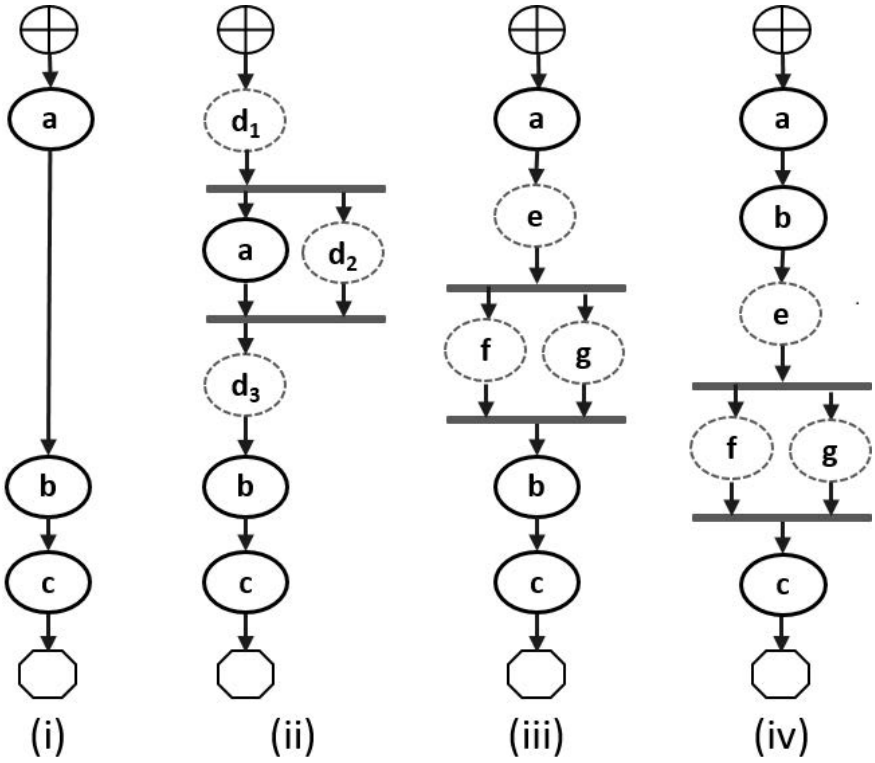


Figure 1—Partial specification of verification intent

Copyright © 2026 Accellera. All rights reserved.

In Figure 1, actions `a`, `b`, and `c` are specified to be traversed sequentially in an activity. Depending on the data flow between them, and on other constraints in the model, this may describe a complete scenario specification (see Figure 1(i)), or it may describe a partial specification, which may be expanded into multiple scenarios that infer other actions. All scenarios satisfy the critical intent defined by the activity, where `a` will be traversed, followed sometime later by `b`, followed sometime later by `c`. Figure 1 shows several possible scenarios that may be generated from the partial specification, depending on various factors to be discussed later in this section.

An *activity* primarily specifies the set of actions to be executed and the scheduling relationships between them. Actions may be scheduled sequentially, in parallel, or in various combinations based on conditional evaluation, looping, or randomization constructs. Activities may also include explicit data bindings between actions. An activity that traverses a compound action is evaluated hierarchically, i.e., when a compound sub-action is traversed in an activity, the sub-action activity is traversed fully at that point in the parent activity (see 5.3.2).

5.1 Modeling data flow

Actions may be declared to have inputs and/or outputs of a given data flow object type. The data flow object types define scheduling semantics for the given action relative to those with which it shares the object. Data flow objects may be declared directly or may inherit from user-defined data structures or other flow objects of a compatible type. An action that outputs a flow object is said to *produce* that object and an action that inputs a flow object is said to *consume* the object. Data flow objects are described in 9.3.

5.1.1 Buffers

The first kind of data flow object is the *buffer* type. A buffer represents *persistent* data that can be written (output) by a producing action and may be read (input) by any number of consuming actions. As such, a buffer defines a strict scheduling dependency between the producer and the consumer that requires the producing action to complete its execution—and, thus, complete writing the buffer object—before execution of the consuming action may begin to read the buffer (see Figure 2). Note that other consuming actions may also input the same buffer object. While there are no implied scheduling constraints between the consuming actions, none of them may start until the producing action completes.

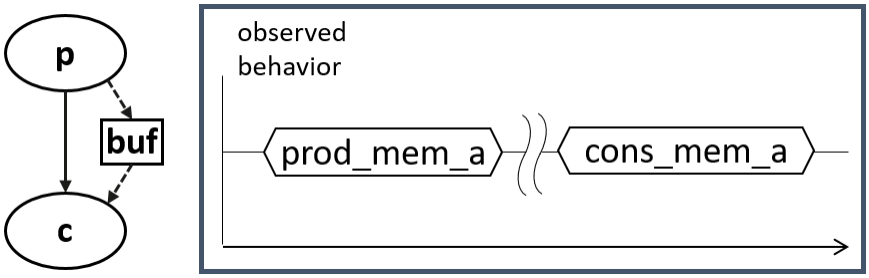


Figure 2—Buffer flow object semantics

Figure 2 illustrates the sequential scheduling semantics between the producer and consumer of a buffer flow object.

In Figure 1(i), assume that action `a` produces a buffer of a particular type, and `b` inputs a buffer object of a compatible type, In this case, we say that the buffer object is *bound* from the output of `a` to the input of `b`, since the semantics of the buffer object support the activity. Similarly, in Figure 1(ii), if, instead of action `a`,

Copyright © 2026 Accellera. All rights reserved.

action d produced a buffer object of a compatible type for action b, action d could be inferred as the producer of the buffer for action b to consume. The buffer scheduling semantics allow action d to be inferred at any point in the schedule prior to the start of action b (shown in Figure 1(ii) as either d_1, d_2, or d_3), while the activity requires only that action a completes before action b starts. In this case, there is no explicit scheduling constraint between a and d.

5.1.2 Streams

The *stream* flow object type represents *transient* data exchanged between actions. The semantics of the stream flow object require that the producing and consuming actions execute in parallel (i.e., both activities shall begin execution when the same preceding actions complete; see Figure 3). In a stream object, there shall be a one-to-one connection between the producer and consumer.

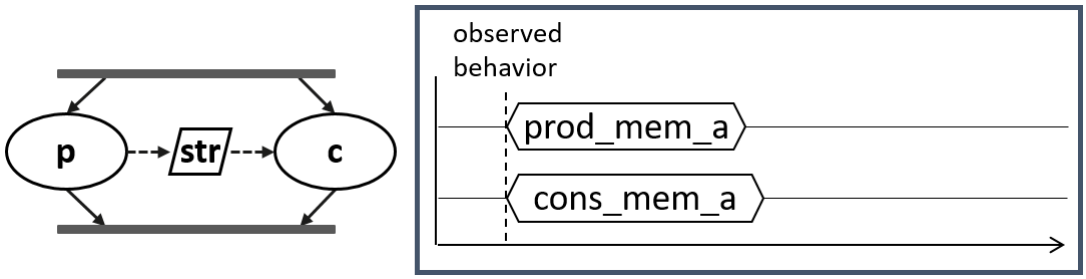


Figure 3—Stream flow object semantics

Figure 3 illustrates the parallel scheduling semantics between the producer and the consumer of a stream flow object.

In Figure 1(iii), the parallel execution of actions f and g dictates that any data exchanged between these actions shall be of the *stream* type. Again, assuming that action a does not output a compatible buffer for action b to input, then action f may be inferred to supply the buffer to action b. If action f inputs or outputs a stream object, then the one-to-one requirement of the stream object would require that action g, which has a compatible stream type, also be inferred to execute in parallel with f. Action e may be inferred if it is needed to supply a buffer input to either f or g.

NOTE—Figure 1(iv) shows an alternate inferred scenario that also satisfies the base scenario of sequential execution of actions a, b, and c, but in this case, the binding between a and b is legal, and action c requires a buffer input that can only be supplied by f or g.

5.1.3 States

The *state* flow object represents the state of some element in the DUT or test environment at a given time. Multiple actions may read or write the state object, but only one write action may execute at a time. Any number of read actions may execute in parallel, but read and write actions shall be sequential (see Figure 4).

Copyright © 2026 Accellera. All rights reserved.

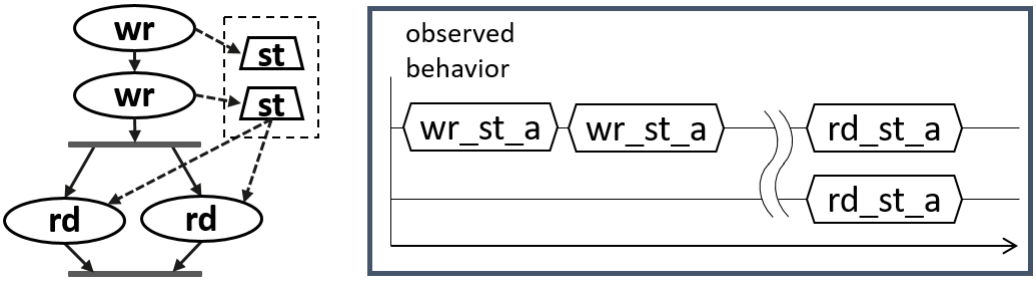


Figure 4—State flow object semantics

State flow objects have a built-in Boolean **`initial`** attribute that is automatically set to **true** initially and automatically set to **false** on the first write operation to the state object. This attribute can be used in constraint expressions to define the starting value for fields of the state object and then allow the values to be modified on subsequent writes of the state object.

5.1.4 Data flow object pools

Data flow objects are grouped into *pools*, which can be used to limit the set of actions that can communicate using objects of a given type. For buffer and stream types, the pool will contain the number of objects of the given type needed to support the communication between actions sharing the pool. For state objects, the pool will only contain a single object of the state type at any given time. Thus, all actions sharing a state object via a pool will see the same value for the state object at a given time. Pools are described in Clause 12.

5.2 Modeling system resources

5.2.1 Resource objects

In addition to declaring inputs and outputs, actions may require system resources that shall be accessible in order to accomplish the specified behavior. The *resource* object is a user-defined data object that represents this functionality. Similar to data flow objects, a resource may be declared directly or may inherit from a user-defined data structure or another resource object. Resource objects are described in 9.4.

5.2.2 Resource pools

Resource objects are also grouped into pools to define the set of actions that have access to the resources. A resource pool is defined to have an explicit number of resource objects in it (the default is `1`), corresponding to the available resources in the DUT and/or test environment. In addition to optionally randomizable data fields, the resource has a built-in non-negative integer attribute called **`instance_id`**, which serves to identify the resource and is unique for each resource in the given pool. Pools are described in Clause 12.

5.2.2.1 Locking resources

An action that requires exclusive access to a resource may *lock* the resource, which prevents any other action that claims the same resource instance from executing until the locking action completes. For a given pool of resource `R`, with size `S`, there may be `S` actions that lock a resource of type `R` executing at any given time. Each action that locks a resource in a given pool at a given time shall have access to a unique instance of the resource, identified by the integer attribute **`instance_id`**. For example, if a DUT contains two DMA channels, the PSS model would define a pool containing two instances of the `DMA_channel` resource type.

Copyright © 2026 Accellera. All rights reserved.

In this case, no more than two actions that lock the `DMA_channel` resource could be scheduled concurrently.

5.2.2.2 Sharing resources

An action that requires non-exclusive access to a resource may *share* the resource. An action may not share a resource instance that is locked by another action, but may share the resource instance with other actions that also share the same resource instance. If all resources in a given pool are locked at a given time, then no sharing actions can execute until at least one locking action completes to free a resource in that pool.

5.3 Basic building blocks

5.3.1 Components and binding

A critical aspect of portability is the ability to encapsulate elements of verification intent into "building blocks" that can be used to combine and compose PSS models. A *component* is a structural element of the PSS model that serves to encapsulate other elements of the model for reuse. A component is typically associated with a structural element of the DUT or testbench environment, such as hardware engines, software packages, or testbench agents, and contains the actions that the element is intended to perform, as well as the data and resource pools associated with those actions. Each component declaration defines a unique type that can be instantiated inside other components. The component declaration also serves as a type namespace in which other types may be declared.

A PSS model is composed of one or more component instantiations constituting a static hierarchy beginning with the top-level or root component, called **`pss_top`** by default, which is implicitly instantiated. Components are identified uniquely by their hierarchical path. In addition to instantiating other components, a component may declare functions and class instances (see Clause 9.1).

When a component instantiates a pool of data flow or resource objects, it also shall *bind* the pool to a set of actions and/or subcomponents to define who has access to the objects in the pool. Actions may only communicate via an object pool with other actions that are bound to the same object pool. Object binding may be specified hierarchically, so a given pool may be shared across subcomponents, allowing actions in different components to communicate with each other via the pool.

5.3.2 Evaluation and inference

A PSS model is evaluated starting with the top-level *root action*, which shall be specified to a tool. The component hierarchy, starting with **`pss_top`** or a user-specified top-level component, provides the context in which the model rules are defined. If the root action is a compound action, its activity forms the root of a potentially hierarchical activity tree that includes all activities present in any sub-activities traversed in the activity. Additional actions may be inferred as necessary to support the data flow and binding requirements of all actions explicitly traversed in the activity, as well as those previously inferred. Resources add an additional set of scheduling constraints that may limit which actions actually get inferred, but resources do not cause additional actions to be inferred.

The semantics of data flow objects allow the tool to infer, for each action in the overall activity, connections to other actions already instantiated in the activity; or to infer and connect new action instances to conform to the scheduling constraints defined in the activity and/or by the data and resource requirements of the actions, including pool bindings. The model thus consists of a set of actions, with defined scheduling dependencies, along with a set of data flow objects that may be explicitly bound or inferred to connect between actions and a set of resources that may be claimed by the actions as each executes. Actions and flow objects and their bindings may only be inferred as required to make the (partial) activity specification legal.

Copyright © 2026 Accellera. All rights reserved.

A PSS implementation shall not infer an action or object binding that is not required, either directly or indirectly, to make the activity specification legal. Clause 14 describes action inferencing in more detail.

Figure 5 demonstrates how actions can be inferred to generate multiple scenarios from a single activity.

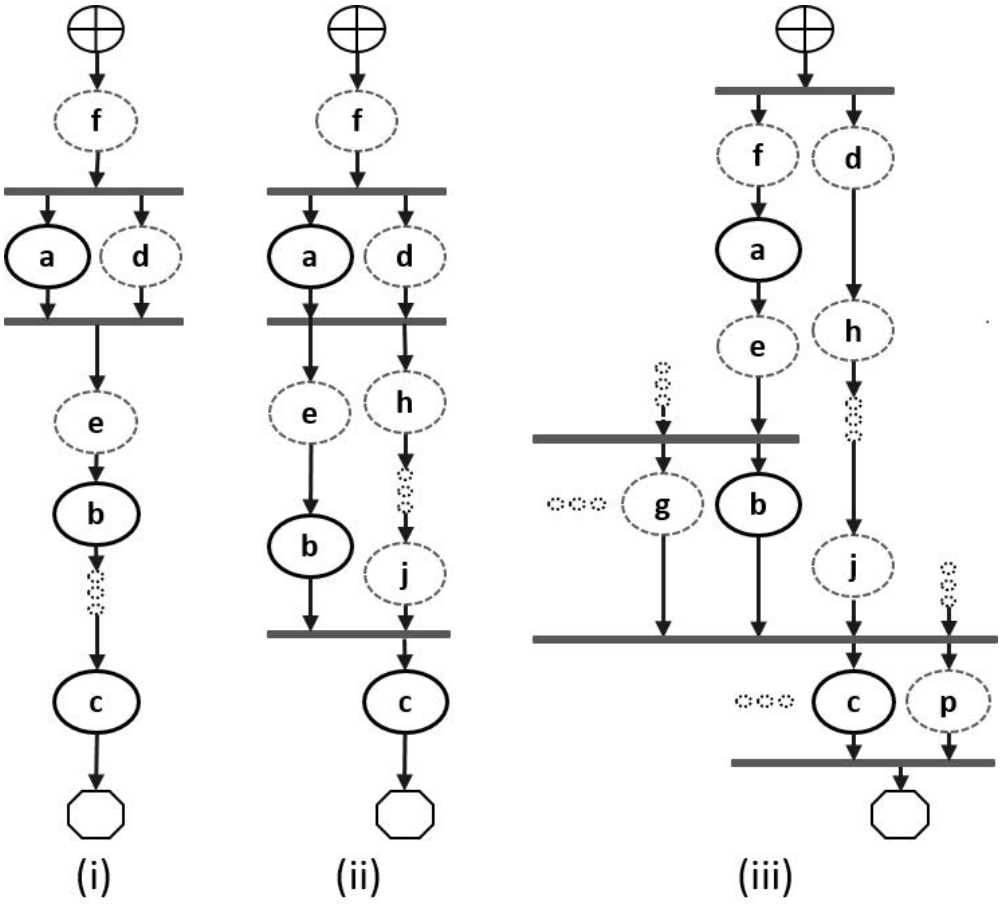


Figure 5—Single activity, multiple scenarios

Looking at Figure 5, actions `a`, `b`, and `c` are scheduled sequentially in an activity. The data flow and resource requirements specified in the model (which are not shown in Figure 5) allow for multiple scenarios to be generated. If action `a` has a buffer or state input, then an action, `f` in this case, is inferred to execute sequentially before `a` in order to provide the buffer or state object. If `a` does not have a buffer or state input, `f` may still be inferred in order to supply an input to `b` or `c`, and may ultimately be scheduled before `a` as shown, although the only real scheduling constraint is that `f` complete before the start of the action that requires the input flow object.

Once inferred, if `f` also has a buffer or state input, then another action shall be inferred to supply that object and so on until an action is inferred that does not have an input (or the tool's inferencing limit is exceeded, at which point an error shall be generated). For the purposes of this example, action `f` does not have an input.

In Figure 5(i), presume that action `a` produces (or consumes) a stream object. In this case, action `d` is inferred in parallel with `a` since stream objects require a one-to-one connection between actions. Actions `a` and `d` both start upon completion of action `f`. If action `d` also has a buffer input, then another action shall be

Copyright © 2026 Accellera. All rights reserved.

inferred to provide that input. For Figure 5(i), action `f` can be presumed to have a second buffer output that gets bound to action `d`, although a second buffer-providing action could also have been inferred.

If action `a` produces a buffer object, the buffer may be connected to another action with a compatible input type. In the case where `a.out` and `b.in` are incompatible, action `e` (or a series of actions) may be inferred to receive the output of action `a` and produce the input to action `b`. If `a.out` and `b.in` are compatible, then the direct connection between `a.out` and `b.in` would be inferred here, in which case no action would be inferred between them, although an action inferred to supply the input to `c` (or for some other reason) could be scheduled between them.

Similarly, in the absence of an explicit binding of `b.out` to `c.in`, and if they are incompatible, a series of actions may be inferred prior to the start of action `c` in order to provide the input of action `c`. These inferred actions will be scheduled independent of `b` unless their data flow requirements create scheduling constraints relative to `b`. As the terminal action in the activity, no action may be inferred after action `c` however, even if action `c` produces a buffer object as an output.

If `b.out` and `c.in` are incompatible, it is possible to infer another action, `j`, to supply the buffer input to `c.in`, as shown in Figure 5(ii). In this case, there are two constraints on when the execution of action `c` may begin. The activity scheduling requires action `b` to complete before action `c` starts. The buffer object semantics also require action `j` to complete before action `c` starts. If action `j` requires a buffer input, a series of actions could be inferred to supply the buffer object. That inferred action chain could eventually be bound to a previously inferred action, such as action `d` as shown in Figure 5(ii), or it may infer an independent series of actions until it infers an initial action that only produces an output or until the inferencing limit is reached. Since the output of action `b` is not bound to action `c`, action `b` is treated as a terminating action, so no subsequent actions may be inferred after action `b`.

Finally, Figure 5(iii) shows the case where action `c` produces or consumes a stream object. In this case, even though action `c` is the terminating action of the activity, action `p` shall be inferred to satisfy the stream object semantics for action `c`. Here, action `p` is also treated as a terminating action, so no subsequent actions may be inferred. However, additional actions may be inferred either preceding or in parallel to action `p` to satisfy its data flow requirements. Each action thus inferred is also treated as a terminating action. Similarly, since action `b` is not bound to action `c`, `b` shall also be treated as a terminating action.

5.4 Constraints and inferencing

Data flow and resource objects may define constraint expressions on the values of their data fields (including **`instance_id`** in the case of resource objects). In addition, actions may also define constraint expressions on the data fields of their input/output flow objects and locked/shared resource objects. For data flow objects, all constraints defined in the object and in all actions that are bound to the object are combined to define the legal set of values available for the object field. Similarly, the constraints defined for a resource object shall be combined with the constraints defined in all actions that claim the resource. Inferred actions or data flow objects that result in constraint contradictions are excluded from the legal scenario. At least one valid solution shall exist for the scenario model for that model to be considered valid.

5.5 Summary

In portable stimulus, a single PSS model may be used to generate a set of scenarios, each of which may have different sets of inferred actions, data flow objects, and resources, while still implementing the critical verification intent explicitly specified in the activity. Each resulting scenario may be generated as a test implementation for the target platform by taking the behavior mapping implementation embedded in each resulting atomic action and generating output code that assembles the implementations and provides any other required infrastructure to ensure the behaviors execute on the target platform according to the scheduling semantics defined by the original PSS model.

Copyright © 2026 Accellera. All rights reserved.

6. Execution semantic concepts

6.1 Overview

A PSS test scenario is identified given a PSS model and an action type designated as the root action. The execution of the scenario consists essentially in executing a set of actions defined in the model, in some (partial) order. In the case of atomic actions, the mapped behavior of any **exec body** clauses (see 20.1.2) is invoked in the target execution environment, while for compound actions the behaviors specified by their **activity** statements are executed.

All action executions observed in a test run either correspond to those explicitly called by traversed activities or are implicitly introduced to establish flows that are correct with respect to the model rules. The order in which actions are executed shall conform to the flow dictated by the activities, starting from the root action, and shall also be correct with respect to the model rules. *Correctness* involves consistent resolution of actions' inputs, outputs, and resource references, as well as satisfaction of scheduling constraints. Action executions themselves shall reflect data attribute assignments that satisfy all constraints.

6.2 Assumptions of abstract scheduling

Guarantees provided by PSS are based on general capabilities that test realizations need to have in any target execution environment. The following are assumptions and invariants from the abstract semantics viewpoint.

6.2.1 Starting and ending action executions

PSS semantics assume that target-mapped behavior associated with atomic actions can be invoked in the execution environment at arbitrary points in time, unless model rules (such as state or data dependencies) restrict doing so. They also assume that target-mapped behavior of actions can be known to have completed.

PSS semantics make no assumptions on the duration of the execution of the behavior. They also make no assumptions on the mechanism by which an implementation would monitor or be notified upon action completion.

6.2.2 Concurrency

PSS semantics assume that actions can be invoked to execute concurrently, under restrictions of model rules (such as resource contentions).

PSS semantics make no assumptions on the actual threading framework employed in the execution environment. In particular, a target may have a native notion of concurrent tasks, as in SystemVerilog simulation; it may provide native asynchronous execution threads and means for synchronizing them, such as embedded code running on multi-core processors; or it may implement time sharing of native execution thread(s) in a preemptive or cooperative threading scheme, as is the case with a runtime operating system kernel. PSS semantics do not distinguish between these.

6.2.3 Synchronized invocation

PSS semantics assume that action invocations can be synchronized, i.e., logically starting at the same time. In practice there may be some delay between the invocations of synchronized actions. However, the "sync-time" overhead is (at worse) relative to the number of actions that are synchronized and is constant with respect to any other properties of the scenario or the duration of any specific action execution.

Copyright © 2026 Accellera. All rights reserved.

PSS semantics make no assumptions on the actual runtime logic that synchronizes native execution threads and put no absolute limit on the "sync-time" of synchronized action invocations.

6.3 Scheduling concepts

PSS execution semantics define the criteria for legal runs of scenarios. The criterion covered in this section is stated in terms of scheduling dependency—the fundamental scheduling relation between action executions. Ultimately, scheduling is observed as the relative order of behaviors in the target environment per the respective mapping of atomic actions. This section defines the basic concepts, leading up to the definition of sequential and parallel scheduling of action executions.

6.3.1 Preliminary definitions

a) An *action execution* of an atomic action type is the execution of its exec-body block,[1] with values assigned to all of its parameters (reachable attributes). The execution of a compound action consists in executing the set of atomic actions it contains, directly or indirectly. For more on execution semantics of compound actions and activities, see Clause 11.

   An atomic action execution has a specific *start-time*—the time in which its exec-body block is entered, and *end-time*—the time in which its exec-body block exits (the test itself does not complete successfully until all actions that have started complete themselves). The start-time of an atomic action execution is assumed to be under the direct control of the PSS implementation. In contrast, the end-time of an atomic action execution, once started, depends on its implementation in the target environment, if any (see 6.2.1).

   The difference between end-time and start-time of an action execution is its *duration*.

b) A *scheduling dependency* is the relation between two action executions, by which one necessarily starts after the other ends. Action execution `b` has a scheduling dependency on `a` if `b`'s start has to wait for `a`'s end. The temporal order between action executions with a scheduling dependency between them shall be guaranteed by the PSS implementation regardless of their actual duration or that of any other action execution in the scenario. Taken as a whole, scheduling dependencies constitute a partial order over action executions, which a PSS solver determines and a PSS scheduler obeys.

   Consequently, the lack of scheduling dependency between two action executions (direct or indirect) means neither one waits for the other. Having no scheduling dependency between two action executions implies that they may (or may not) overlap in time.

c) Action executions are *synchronized* (scheduled to start at the same time) if they all have the exact same scheduling dependencies. No delay shall be introduced between their invocations, except a minimal constant delay (see 6.2.3).

d) Two or more sets of action executions are *independent* (scheduling-wise) if there is no scheduling dependency between any two action executions across the sets. Note that within each set, there may be scheduling dependencies.

e) Within a set of action executions, the *initial* ones are those without scheduling dependency on any other action execution in the set. The *final* action executions within the set are those in which no other action execution within the set depends.

[1] Throughout this section, exec-body block is referred to in the singular, although it may be the aggregate of multiple exec-body clauses in different locations in PSS source code (e.g., multiple declarations in a given action type definition or in different extensions of the same action type).

Copyright © 2026 Accellera. All rights reserved.

6.3.2 Sequential scheduling

Action executions `a` and `b` are scheduled in *sequence* if `b` has a scheduling dependency on `a`. Two sets of action executions, S_1 and S_2, are scheduled in sequence if every initial action execution in S_2 has a scheduling dependency on every final action execution in S_1. Generally, sequential scheduling of N action execution sets S_1 .. S_n is the scheduling dependency of every initial action execution in S_i on every final action execution in S_{i-1} for every i from 2 to N, inclusive.

For examples of sequential scheduling, see 11.3.3.2.

6.3.3 Parallel scheduling

N sets of action executions S_1 .. S_n are scheduled in *parallel* if the following two conditions hold:

— All initial action executions in all `N` sets are synchronized (i.e., all have the exact same set of scheduling dependencies).

— S_1 .. S_n are all scheduled independently with respect to one another (i.e., there are no scheduling dependencies across any two sets S_i and S_j).

For examples of parallel scheduling, see 11.3.4.2.

6.3.4 Concurrent scheduling

N sets of action executions S_1 .. S_n are scheduled *concurrently* if S_1 .. S_n are all scheduled independently with respect to one another (i.e., there are no scheduling dependencies across any two sets S_i and S_j).

6.4 Test realization scheduling assumptions

Test realization for atomic *actions* within a scenario is provided by *exec blocks*. During test execution, exec blocks are mapped to one or more hardware threads (*executors*).

The following preliminary definitions apply:

a) Test realization exec blocks are mapped to one or more *executors* (hardware threads) during test execution. See section 21.7 for more details on executors.

b) Each *executor* represents a single thread of execution. Executors should have the ability to run multiple exec blocks that are simultaneously executing on the same executor using techniques like cooperative or preemptive multitasking. See 20.8 for more details on test realization cooperative multitasking.

Copyright © 2026 Accellera. All rights reserved.

7. Data types

7.1 General

In this document, *numeric types* are the data types **bit**, **int**, **float32**, or **float64**, as well as any derivatives of **bit** and **int**. A single data item of a numeric type is a *"number"*. "Scalar types" are the numeric types, **bool**, **string**, **chandle**, and all enumeration types. A single data item of a scalar type is a "*scalar*". A **struct** (see 7.8) or *collection* (see 7.9) is not a scalar. A **typedef** (see 7.11) of a scalar data type is also a scalar data type. A field of plain-data type may be declared as constant by preceding its declaration with the **const** keyword. If the constant is of aggregate type, its elements are constants too. The value of constant fields can be read but not modified. A constant cannot appear on the left-hand side of the assignment operator.

The term "*aggregate"* refers both to *collections* and to **structs**. The term "*aggregate*" does not include **actions**, **components**, **monitors**, *flow objects*, or *resource objects*, also known as modeling elements (see Clause 9). Aggregates may be nested. A **typedef** of an aggregate data type is also an aggregate data type.

A "*plain-data type*" is a scalar or an aggregate of scalars. Nested aggregates are also plain-data types. A **typedef** of a plain-data type is also a plain-data type.

Fields of all scalar types except **chandle**, **float32**, and **float64** are *randomizable*. Array and list collections of randomizable types are also randomizable, but the **map** and **set** collection types are not randomizable.

A field of randomizable type may be declared as *random* by preceding its declaration with the **rand** keyword. It shall be an error to declare a field of non-randomizable type as **rand**.

7.1.1 Syntax

The syntax for data types and data declarations is shown in Syntax 8.

```
data_type ::=
    scalar_data_type
  | collection_type
  | reference_type
  | type_identifier
scalar_data_type ::=
    chandle_type
  | integer_type
  | string_type
  | bool_type
  | enum_type
  | float_type
data_declaration ::= data_type data_instantiation { , data_instantiation } ;
data_instantiation ::= identifier { array_dim } [ = constant_expression ]
array_dim ::= [ constant_expression ]
attr_field ::= [ access_modifier ] [ rand | static const ] data_declaration
access_modifier ::= public | protected | private
```

Syntax 8—Data types and data declarations

Copyright © 2026 Accellera. All rights reserved.

Scalar data types are described in 7.2 through 7.7, structure data types are described in 7.8, and collection data types are described in 7.9. Reference types are described in 7.10. Access protection and access modifiers are described in 17.4.

7.2 Integer types

PSS supports two 2-state integer data types. These fundamental integer data types are summarized in Table 6, along with their default widths and value domains.

Table 6—Integer data types

Data type	Default width	Default domain	Signed/Unsigned
int	32 bits	-2^{31} .. $(2^{31}-1)$	Signed
bit	1 bit	0..1	Unsigned

4-state values are not supported. If 4-state values are passed into the PSS model via the *foreign procedural interface* (see 20.4), any **x** or **z** values are converted to **0**.

7.2.1 Syntax

The syntax for integer types is shown in Syntax 9.

```
integer_type ::= integer_atom_type
   [ [ constant_expression ] ]
   [ in [ domain_open_range_list ] ]
integer_atom_type ::=
    int
  | bit
domain_open_range_list ::= domain_open_range_value { , domain_open_range_value }
domain_open_range_value ::=
    constant_expression [ .. constant_expression ]
  | constant_expression ..
  | .. constant_expression
```

Syntax 9—Integer type declaration

The following also apply:

a) Integer values of **bit** type are unsigned. Integer values of **int** type are signed.

b) The default value of the integer types is 0, with the exception shown in clause d below.

c) Widths should be specified with a single expression with a constant positive integer value (e.g., `bit[4]`).

d) A value domain may be specified for the type. The domain specification consists of a list of one or more values and/or value ranges. The values and range bounds shall be evaluated as self-determined expressions (see 8.7.1). The default value of a type whose value domain does not include 0 is the

Copyright © 2026 Accellera. All rights reserved.

smallest value in the domain. When a variable of a type with a restricted domain is randomized, the domain shall serve as an implicit set membership constraint (see 8.5.9 and 13.1.5).

e) The width and value domain specifications are independent. A variable of the declared type can hold values within the intersection of the possible values determined by the specified width (or the default width, if not specified) and the explicit value domain specification, if present.

7.2.2 Examples

PSS integer data type examples are shown in-line in this section.

Declare a signed variable that is 32 bits wide.

```
int a;
```

Declare a signed variable that is 5 bits wide.

```
int [5] b;
```

Declare an unsigned variable that is 5 bits wide and has the valid values `1..31`.

```
bit [5] in [1..31] c;
```

Declare an unsigned variable that is 5 bits wide and has the valid values `1`, `2`, and `4`. The default value is 1.

```
bit [5] in [1,2,4] d;
```

Declare an unsigned variable that is 5 bits wide and has the valid values `0..10`.

```
bit [5] in [..10] e; // 0 <= e <= 10
```

Declare an unsigned variable that is 5 bits wide and has the valid values `10..31`. The default value is 10.

```
bit [5] in [10..] f; // 10 <= f <= 31
```

Declare a signed variable that is 32 bits wide and has the valid values -10..5. The default value is 0.

```
int in [-10..5] g;
```

Declare a signed variable that is 32 bits wide and has the valid values -10..-5. The default value is -10.

```
int in [-10..-5] h;
```

Declare a signed variable that is 5 bits wide and has the valid values -16..-5. The default value is -16.

```
int [5] in [..-5] i; // -16 <= i <= -5
```

7.3 Floating-point types

PSS supports two floating-point *computation* data types, as summarized by Table 7 below.

7.3.1 Syntax

The syntax for floating-point computation data types is shown in Syntax 10 below.

Copyright © 2026 Accellera. All rights reserved.

Table 7—Floating-point computation data types

Data type	Width	Format
float32	32 bits	IEEE 754 binary32
float64	64 bits	IEEE 754 binary64

```
scalar_data_type ::=
    ...
  | float_type
float_type ::=
    float32
  | float64
```

Syntax 10—Floating-point type declaration

Variables of floating-point type may not be declared **rand**, and may not be randomized using the **randomize** statement.

PSS also defines packed-struct *storage* types as part of the core library (see 21.13.1). These types support various non-IEEE floating-point number formats.

Arithmetic operations may be performed on the *computation* data types. Arithmetic operations may not be performed directly on storage data types. Data held in a variable of floating-point storage type shall first be converted into a computation type.

7.3.2 Cross-platform results

Floating-point computation has platform dependencies, with different processors and algorithms legitimately producing slightly different results. These differences may be apparent, for example, when comparing the result of computations performed on the solve platform with those performed on the target platform. The PSS LRM makes no attempt to force the result of floating-point computations to be identical across platforms.

7.4 Booleans

The PSS language supports a built-in Boolean type, with the type name **bool**. The **bool** type has two enumerated values **true** (=`1`) and **false** (=`0`). When not initialized, the default value of a **bool** type is **false**.

7.5 Enumeration types

An *enumeration type* is a distinct user-defined type whose value is restricted to a specified set of integral named constants. Enumeration data types also can be easily referenced or displayed using the enumeration constant names as opposed to their numeric values.

7.5.1 Syntax

The syntax for declaration of enumeration types is shown in Syntax 11.

Copyright © 2026 Accellera. All rights reserved.

```
enum_declaration ::= enum enum_identifier [ : data_type ] { [ enum_item { , enum_item } ] }
enum_identifier ::= identifier
enum_item ::= identifier [ = constant_expression ]
enum_type_identifier ::= type_identifier
enum_type ::= enum_type_identifier [ in [ domain_open_range_list ] ]
```

Syntax 11—Enum declaration

An enumeration type declaration (*enum_declaration*) consists of the keyword **enum** followed by the name of the type (*enum_identifier*), an optional base type name (*data_type*), and a list in curly braces of constant names (*enum items*) with optional constant integer value assignments.

The optional *data_type* denotes the base type. It shall be the name of an integer type, which shall determine the set of possible values to be assigned to *enum_item*s, for example: **int**, or **bit**[16], or **int**[3]. In effect, it shall determine the width and the signedness of the items. The base type shall not have a value domain (for example, '**int in** [1..10]' cannot be used as a base type).

The following also apply:

a) *enum_item*s are considered static constant members of the enumeration type in which they are declared.

b) The first *enum_item* in the list, if not explicitly assigned a value, is by default assigned the value `0`. Each following *enum_item*, if not explicitly assigned a value, is assigned a value of the previous *enum_item* `+ 1`.

c) If a base type (*data_type*) is specified, *enum_item* values are limited to the set of valid values of the base type. It shall be an error to explicitly assign a value which does not belong to the base type (for example, if the base type is unsigned, it shall be an error to assign a negative value). It shall also be an error to declare an *enum_item* without an explicit value if the previous *enum_item* has been assigned the greatest possible value of the base type (for example, if the base type is **bit**[2], declaring an item without an explicit value is illegal if the previous item has the value 3).

d) *enum_item* values need not be contiguous, nor need they be in ascending arithmetic order. An *enum_item* may be assigned a negative value (unless the base type is unsigned).

e) Each *enum_item* shall have a distinct integer value. No two *enum_items* may have the same value.

f) Enumeration types may be *extended* with the **extend** statement. See 17.2, particularly 17.2.4.

g) *enum_item* identifiers shall be unique in the scope of the enumeration type across its initial definition and extensions, if any. However, they need not be unique across different enumeration types declared in the same namespace.

h) *enum_items* can be referenced using their *qualified name* in the form `'enum-type-name::enum-item-name'`.

i) In expression contexts where the expected type is an enumeration type, *enum_items* of that type can be referenced without qualification (see 8.4.3 for the definition of the expected type in expression contexts).

j) An *enum_declaration* may contain an empty set of *enum_item*s, and then have *enum_item*s added in extensions. It shall be illegal to declare an enumeration variable whose type contains no *enum_items* across its initial definition and extensions.

k) When not initialized, the default value of an **enum** field shall be the first *enum_item* in the list. This is not necessarily the value `0` nor the *enum_item* with the minimum value.

Like numeric types, an enumeration type can be restricted to a range of values specified by a *domain_open_range_list* (see 7.2.1 and 7.2.2). The domain specification cannot be specified in the

Copyright © 2026 Accellera. All rights reserved.

enum_declaration itself. See examples of use in 7.5.2. The default value of a type with a restricted domain is the first declared *enum_item* among those that belong to the domain. When a variable of a type with a restricted domain is randomized, the domain shall serve as an implicit set membership constraint (see 8.5.9 and 13.1.5).

An **enum** attribute or *enum_item* may be used to assign values to an attribute of the same enumeration type or in an equality comparison.

An **enum** attribute or *enum_item* of one enumeration type may be cast to another enumeration type using the cast operator (see 7.12). An **enum** attribute or *enum_item* may be cast to integer and Boolean data types using the cast operator. Similarly, an integer or Boolean value may be explicitly cast to an enumeration type.

7.5.2 Examples

Examples of enum usage are shown in Example 8.

```
enum config_modes_e {UNKNOWN, MODE_A=10, MODE_B=20, MODE_C=35, MODE_D=40};

component uart_c {
  action configure {
    rand config_modes_e mode;
    constraint {mode != UNKNOWN;}
  }
};
```

Example 8—Enum data type

See an example of extending an enumeration in 17.2.4.

Examples of domain specifications for enumeration types are shown below:

Declare an enum of type `config_modes_e` with values `MODE_A`, `MODE_B`, or `MODE_C`.

```
rand config_modes_e in [MODE_A..MODE_C] mode_ac;
```

Declare an enum of type `config_modes_e` with values `MODE_A` or `MODE_C`. The default value is `MODE_A`.

```
config_modes_e in [MODE_A, MODE_C] mode_a_c;
```

Declare an enum of type `config_modes_e` with values `UNKNOWN`, `MODE_A`, or `MODE_B`.

```
rand config_modes_e in [..MODE_B] mode_ub;
```

Declare an enum of type `config_modes_e` with values `MODE_B`, `MODE_C`, or `MODE_D`.

```
rand config_modes_e in [MODE_B..] mode_bd;
```

Note that an *open_range_list* of enums may be used in set membership (**in**) expressions (see 8.5.9) and as a *match_choice* expression in **match** statements (see 11.4.6 and 20.7.10).

Copyright © 2026 Accellera. All rights reserved.

7.6 Strings

The PSS language supports a built-in string type with the type name **string**. When not initialized, the default value of a **string** shall be the empty string literal (**""**).

7.6.1 Syntax

```
string_type ::= string [ in [ string_literal { , string_literal } ] ]
```

Syntax 12—String declaration

Comma-separated domain specifications are allowed for string data types (see 7.2.1). The following applies to the sub-string operator (see 7.6.2), all string methods (see 7.6.3), and string-related collection methods *join()* and *str_from_chars()* (see 7.9.2.2 and 7.9.3.2): in some environments (for example, certain embedded-software environments), the usage of these operators and methods may be limited in context of target execs if the values of all parameters are not known at solve time. This is due to the target platform memory requirements for the string operations or other considerations.

7.6.2 The sub-string operator

The *sub-string operator* is used to get a sub-string from a given string, given starting and/or ending character positions within the string. See 8.6.3 for more information on the sub-string operator.

The sub-string operator shall not be used on the left-hand side of an assignment operator.

The sub-string operator is not randomizable. Therefore, it can be used in constraints only if neither the string nor the slice indices are themselves randomized.

7.6.3 String methods

The following methods are defined for **string**s. In all methods, character positions are counted starting from 0, i.e., the position of the first character in the string is 0.

pure function int *size();*

Returns the size of the string, i.e., the number of characters it contains.

pure function int *find*(**string** *sub_str,* **int** *first_pos* = 0);

Returns the starting position of *sub_str* within the string, counting from *first_pos* (0 by default). If *sub_str* appears within the string more than once, the starting position of the first occurrence is returned. If *sub_str* is not found within *str*, -1 is returned. If *sub_str* is an empty string, 0 is returned, regardless of the string value. Valid values of *first_pos* are between 0 and the string size.

pure function int *find_last*(**string** *sub_str,* **int** *first_pos* = -1);

Returns the starting position of the last occurrence of *sub_str* within the string, counting from *first_pos* backwards. If *first_pos* is -1 (the default), the backwards search is done from the end of the string. If *sub_str* is not found within *str*, -1 is returned. If *sub_str* is an empty string, the size of the string is returned. Valid values of *first_pos* are between -1 and the string size.

Copyright © 2026 Accellera. All rights reserved.

pure function list<int> *find_all*(**string** *sub_str*);

Returns the list of starting positions of all occurrences of *sub_str* within the string, in increasing order. If *sub_str* is not found within *str*, an empty list is returned. If *sub_str* is an empty string, a list of all numbers between 0 and the size of the string is returned.

pure function string *lower*();

Converts all upper-case letters in the string to lower case, and returns the resulting string.

pure function string *upper*();

Converts all lower-case letters in the string to upper case, and returns the resulting string.

pure function list<string> *split*(**string** *sep*);

Splits the string by the separator string *sep*, and returns a list of strings containing the separated sub-strings.

If *sep* occurs at the very beginning or end of the string, the resulting list has an empty string as the first or last element, respectively. If two or more occurrences of *sep* within the string are adjacent, the relevant elements of the resulting list are empty strings.

If the string is empty, the resulting list contains one element, which is an empty string, regardless of the value of *sep*.

sep shall not be an empty string.

pure function list<bit[8**]>** *chars*();

Returns the list of 8-bit character (ASCII) codes of the characters in the string. If the string is empty, an empty list is returned.

See also array and list methods *join*() and *str_from_chars*() in 7.9.2.2 and 7.9.3.2.

7.6.4 Example

The value of a random string-type field can be constrained with equality constraints and can be compared using equality operators, as shown in Example 9.

```
struct string_s {
  rand bit      a;
  rand string   s;

  constraint {
    if (a == 1) {
      s == "FOO";
    } else {
      s == "BAR";
    }
  }
}
```

Example 9—String data type

Declare string with values `"Hello"`, `"Hallo"`, or `"Ni Hao"`.

Copyright © 2026 Accellera. All rights reserved.

```
rand string in ["Hello", "Hallo", "Ni Hao"] hello_s;
```

Note that an *open_range_list*, composed solely of individual string literals, may also be used in set membership (**in**) expressions (see 8.5.9) and as a *match_choice* expression in **match** statements (see 11.4.6 and 20.7.10). Ranges of string literals (e.g., `"a".."b"`) are not permitted.

Copyright © 2026 Accellera. All rights reserved.

Example 10 shows the use of string operators and methods.

```
component pss_top {
  action string_manipulations {
    string hello_str = "Hello, world, and good bye!";
    int n=1, m=7, k=11;

    string s1, s2, s3, s4, s5, s6, s7, s8, s9, s10, s11;
    int str_size, n1, n2, n3, n4, n5;
    list<int> int_l;
    list<string> str_l1, str_l2, str_l3;
    list<bit[8]> char_codes1, char_codes2;

    exec post_solve {
      s1 = hello_str[..4];       // s1 = "Hello"
      s2 = hello_str[m..k];      // s2 = "world"
      s3 = hello_str[m..];       // s3 = "world, and good bye!"
      s4 = hello_str[1];         // s4 = "e"
      // s5 = hello_str[7..100]; // ERROR: too large ending position
      // s6 = hello_str[4..1];   // ERROR: starting larger than ending

      str_size = hello_str.size();  // str_size = 27


      n1 = hello_str.find("world");    // n1 = 7
      n2 = hello_str.find("earth");    // n2 = -1
      n3 = hello_str.find("o");        // n3 = 4
      n4 = hello_str.find("o", 6);     // n4 = 8
      n5 = hello_str.find_last("o");   // n5 = 20
      int_l = hello_str.find_all("o"); // int_l = {4,8,19,20}

      s7 = hello_str.lower(); // s7 = "hello, world, and good bye!"
      s8 = hello_str.upper(); // s8 = "HELLO, WORLD, AND GOOD BYE!"

      str_l1 = hello_str.split(", ");
            // str_l1 = {"Hello", "world", "and good bye!"}
      str_l2 = "abc123abcabc456".split("abc");
            // str_l2 = {"", "123", "", "456"}

      str_l3 = {"ABC", "XYZ", "123"};
      s9 = str_l3.join("::");  // s9 = "ABC::XYZ::123"
      s10 = str_l3.join("");   // s10 = "ABCXYZ123"

      char_codes1 = s1.chars(); // "Hello" -> {72,101,108,108,111}
      char_codes2 = {65,66,67,68,69};
      s11 = char_codes2.str_from_chars(); // s11 = "ABCDE"
    }
  }
}
```

Example 10—String operators and methods

7.7 Chandles

The **chandle** type (pronounced “see-handle”) represents an opaque handle to a foreign language pointer as shown in Syntax 13. A **chandle** is used with the foreign procedural interface (see 20.4) to store foreign

Copyright © 2026 Accellera. All rights reserved.

language pointers in the PSS model and pass them to foreign language functions. See Annex D for more information about the foreign procedural interface.

A **chandle** has the following restrictions:

— The **rand** qualifier may not be applied to it.
— The only logical operators it may be used with are **==** and **!=**.
— The only literal value with which it may be compared is **0**, which is equivalent to a null handle in the foreign language.

When not initialized, the default value of a **chandle** shall be **0**.

7.7.1 Syntax

```
chandle_type ::= chandle
```

Syntax 13—Chandle declaration

7.7.2 Example

Example 11 shows a **struct** containing a **chandle** field that is initialized by the return of a foreign language function.

```
function chandle do_init();

struct info_s {
  chandle ptr;

  exec pre_solve {
    ptr = do_init();
  }
}
```

Example 11—Chandle data type

Copyright © 2026 Accellera. All rights reserved.

7.8 Structs

A **struct** type is an aggregate of data items, as shown in Syntax 14.

7.8.1 Syntax

```
struct_declaration ::= struct_kind struct_identifier [ template_param_decl_list ]
    [ struct_super_spec ] { { struct_body_item } }
struct_kind ::=
    struct
  | object_kind
object_kind ::=
    buffer
  | stream
  | state
  | resource
struct_super_spec ::= : type_identifier
struct_body_item ::=
    constraint_declaration
  | generic_constraint_declaration
  | attr_field
  | typedef_declaration
  | exec_block_stmt
  | attr_group
  | compile_assert_stmt
  | covergroup_declaration
  | covergroup_instantiation
  | struct_body_compile_if
  | annotation_usage struct_body_item
  | stmt_terminator
```

Syntax 14—Struct declaration

A **struct** is a plain-data type (see 7.1). That is, a **struct** may contain scalar data items and aggregates thereof. A **struct** declaration may specify a *struct_super_spec*, a previously defined **struct** type from which the new type inherits its members, by using a colon (**:**), as in C++. In addition, **struct**s may

- include **constraint**s (see 13.1) and **covergroup**s (see 15.1 and 15.2);
- include **exec** blocks of any kind other than **init_down**, **init_up**, and **body** (see 20.1)**.**

Data items in a **struct** shall be of plain-data types (whether randomizable or not). Declarations of randomizable data items may optionally include the **rand** keyword to indicate that the element shall be randomized when the overall **struct** is randomized (see Example 12). 13.4.1 describes **struct** randomization in detail.

Copyright © 2026 Accellera. All rights reserved.

7.8.2 Example

A **struct** example is shown in Example 12.

```
struct axi4_trans_req {
    rand bit[32]     axi_addr;
    rand bit[32]     axi_write_data;
    bit              is_write;
    rand bit[4]      prot;
    rand bit[2]      sema4;
}
```

Example 12—Struct with rand qualifiers

7.9 Collections

Collection types are built-in data types. PSS supports fixed-size **array** and variable-size **list**, **map**, and **set** collections of plain-data types (see 7.1). Each kind of collection has its own keyword, and its declaration specifies the data type of the collection elements (and for **map**s, also the data type of the *key*).

PSS also has limited support for fixed-sized arrays of modeling elements, as described in 7.9.2. These are not considered plain-data types. All other collections are plain-data types.

7.9.1 Syntax

```
collection_type ::=
    array < data_type , array_size_expression >
  | list < data_type >
  | map < data_type , data_type >
  | set < data_type >
array_size_expression ::= constant_expression
```

Syntax 15—Collection data types

In an **array**, each element is initialized to the default initial value of the element type, unless the **array** declaration contains an initialization assignment. A **list**, **map** or **set** is initialized as an empty collection unless the declaration contains an initialization assignment. A collection that is empty is as if it was assigned an *empty aggregate literal* ({}). See 4.8 for more information on literal syntax and semantics used to initialize collection types.

Collections store both scalar and aggregate elements by value. This means that an element's value is captured when it is added or assigned to a collection. Modifying the value of an element in a collection does not modify the element originally added to the collection. In the example below, `v1`, a **struct** with two integer values, is assigned as the first element of `my_list`. Modifying `a` in that element does not modify `v1`. (See 7.9.3 for more details on **list** operators and methods.)

Copyright © 2026 Accellera. All rights reserved.

```
struct my_s1 {
  int a, b;
}

struct my_s2 {
  list<my_s1> my_list;

  exec pre_solve {
    my_s1 v1 = {.a=1,.b=2};
    my_list.push_back(v1);
    my_list[0].a = 10; // my_list == {{.a=10,.b=2}}, v1 == {.a=1,.b=2}
  }
}
```

Example 13—Modifying collection contents

Collection variables can be operated on with built-in operators using standard operator symbols (e.g., **[],** =, ==, etc.) or with built-in methods using a method name and an argument list in parentheses.

Operators and methods that modify the contents of a collection shall not be used in activities, constraints, or **covergroup**s. These are allowed only in **exec** blocks (see 20.1) and native functions (see 20.3). Operators and methods that do not modify collection contents may be used in activities, constraints, and **covergroup**s.

Operators and methods that modify the contents of a collection shall not be used on collections declared with the *const* qualifier. The following also apply on constant collections:

a) The *Index operator* `[]` returns constant.

b) They cannot appear on the left-hand side of the *assignment operator*.

c) The iterator variable of the *foreach statement* will be considered constant inside the loop.

Arrays and **list**s of randomizable types are randomizable. **Map**s and **set**s are non-randomizable. It is legal to have a **rand** struct field that contains non-randomizable collection types.

Collection types may be nested to describe more complex collections.

```
struct my_s {
  list<map<string, int>> m_list_of_maps;
  map<string, list<int>> m_map_of_lists;
}
```

Example 14—Nested collection types

7.9.2 Arrays

PSS supports fixed-sized arrays of plain-data types. Arrays may be declared with two different syntaxes, the classical syntax where arrays are declared by adding square brackets with the array size (**[** *constant_expression* **]**) after the array name, referred to as the *square array* syntax, and the syntax that is aligned to the other collection types, using angle brackets, referred to as the *template array* syntax.

Copyright © 2026 Accellera. All rights reserved.

```
int my_int_arr1[20];              // Square array declaration syntax
array<int,20> my_int_arr2;        // Template array declaration syntax

// Single-dimensional array
int my_int_arr1[20];              // Square array declaration syntax
array<int,20> my_int_arr2;        // Template array declaration syntax

// Two-dimensional array
int my_2d_int_arr1[3][4];              // Square array declaration syntax
array<array<int,4>,3> my_2d_int_arr2; // Template array declaration syntax
```

Example 15—Array declarations

The same operators and methods may be applied to arrays declared using both syntaxes. However, the template array syntax may be used where a *data_type* is required, enabling such capabilities as use as a function return type, arrays nested within other collections, and more.

An array with `N` elements, is ordered, with the first element accessed using **0** as an index value with the **[]** operator, and the last element accessed using `N-1` as an index value.

The square array syntax can also be used to declare fixed-size arrays, possibly multidimensional, of *modeling elements*. Individual elements of such arrays may be accessed using the **[]** operator. However, other operators and methods do not apply to these arrays, unless otherwise specified. Action handle arrays are described in 11.3.1.1 and 11.3.2, monitor handle arrays are described in 16.3.9, component arrays are described in 9.1.4, and object reference arrays are described in 9.3.4 and 9.4.2. Note that the elements of action and monitor handle arrays and object reference arrays have reference semantics (see 7.10).

The template array syntax cannot be used for components, actions, monitors, flow and resource objects, with one exception: it can be used to declare a flow or resource object array as a data field within another flow or resource object (see 9.3.1.1 and 9.4.1.1). However, this syntax can be used with **ref** types (see 7.10).

7.9.2.1 Array operators

The following operators are defined for **array**s:

Index operator **[]**

Used to access a specific element of an array, given an index into the array. The index shall be an integral value. See 8.6.2 for more information on the index operator.

Assignment operator **=**

Creates a copy of the array-type expression on the RHS and assigns it to the array on the LHS. See 8.3 for more information on the assignment operator.

Equality operator **==**

Evaluates to *true* if all elements with corresponding indexes are equal. Two arrays of different element types or different sizes are incomparable. See 8.5.3 for more information on the equality operator.

Inequality operator **!=**

Evaluates to *true* if not all elements with corresponding indexes are equal. Two arrays of different element types or different sizes are incomparable. See 8.5.3 for more information on the inequality operator.

Copyright © 2026 Accellera. All rights reserved.

Set membership operator **in**

The set membership operator can be applied to an array to check whether a specific element is currently within the array. It evaluates to *true* if the element specified on the left of the operator exists in the **array** collection on the right of the operator. The type of the element shall be the same as the array's element data type. See 8.5.9 for more information on the set membership operator.

foreach *statement*

The **foreach** statement can be applied to an array to iterate over the array elements within an activity, a constraint or native exec code. See 11.4.3, 13.1.8, and 20.7.8, respectively, for more information on the **foreach** statements in these contexts.

7.9.2.2 Array methods

The following methods are defined for **array**s:

pure function int *size()***;**

Returns the number of elements in the array. Since arrays have fixed sizes, the returned value is considered a constant expression. This function can also be used with arrays of *modeling elements*.

pure function <*data_type*> *sum()***;**

Returns the sum of all elements currently stored in the array. This function can only be used on arrays of a numeric data type (**int**, **bit**, or floating-point type). The method can be used in a constraint to constrain an array of random **int** or **bit** elements to have a sum of a certain value.

The return type of this function is dependent on the type of the data element:

Table 8—Return type of sum() function

Data type	Return type
int, bit	**int**
float32, float64	**float64**
Other (e.g., **string**, **struct**)	Not applicable

pure function string *join*(**string** *connector*);

This method only applies to arrays of **strings**. It concatenates the strings from the array using *connector* as the connector between them and returns the resulting string.

If *connector* is an empty string, *join*() concatenates the strings with no whitespace inserted.

If the array size is 1, the string value of this element itself is returned and *connector* is ignored.

pure function string *str_from_chars*();

Copyright © 2026 Accellera. All rights reserved.

This method only applies to arrays of integer types. It returns the string whose characters' (ASCII) codes appear in the array, in the same order. Depending on the specific type of the array elements, each value is implicitly cast to 8 bits that represent the character code.

pure function list<*data_type*> *to_list()***;**

Returns a **list** containing the elements of the array. The **list**'s element data type is the same as the data type of the array elements. The **list** elements are ordered in the same order as the array.

pure function set<*data_type*> *to_set()***;**

Returns a **set** containing the elements of the array. Each element value will appear once. The **set**'s element data type is the same as the data type of the array elements. The **set** is unordered.

7.9.2.3 Examples

Examples of fixed-size array declarations are shown in Example 16.

```
int fixed_sized_arr [16];          // array of 16 signed integers
array<bit[8],256> byte_arr;      // array of 256 bytes
array<route,8>      east_routes;  // array of 8 route structs
```

Example 16—Fixed-size arrays

In Example 16, individual elements of the `east_routes` array are accessed using the index operator **[]**, i.e., `east_routes[0]`, `east_routes[1]`,….

Example 17 shows use of array operators and methods. In this example, action type `A` is traversed six times, once for each element in `foo_arr`, and once more since `foo_arr[0]` is greater than `3`.

Copyright © 2026 Accellera. All rights reserved.

```
component pss_top {
  array<bit[16],5> foo_arr;
  set  <bit[16]>   foo_set;

  exec init_up {
    foo_arr = {1, 2, 3, 4, 4};    // Array initialization assignment
    foo_arr[0] = 5;               // Use of [] to select an array element
    foo_set = foo_arr.to_set();   // Use of to_set() method
  }

  action A{ rand bit[16] x; }
  action B{}
  action C{}

  action traverse_array_a {

    // foo_arr has 5 elements and foo_set has 4
    rand int in [1..] y;
    constraint y < comp.foo_arr.size(); // Use of size() method in constraint

    activity {
      foreach (elem: comp.foo_arr)           // "foreach" used on an array
        do A with { x == elem; };

        if (comp.foo_arr[0] > 3)
          do A;
        else if (4 in comp.foo_arr)          // Use of "in" operator
          do B;
        else if (comp.foo_arr.size() < 4)    // Use of size() method
          do C;
    }
  }
}
```

Example 17—Array operators and methods

7.9.3 Lists

The **list** collection type is used to declare a variable-sized ordered list of elements. Using an index, an element in the list can be assigned or used in an expression. A list with `N` elements, is ordered, with the first element accessed using **0** as an index value with the **[]** operator, and the last element accessed using `N-1` as an index value.

A **list** is initialized as an empty collection unless the declaration contains an initialization assignment. A **list** that is empty is as if it was assigned an *empty aggregate literal* (**{}**). **List** elements can be added or removed in **exec** blocks; therefore the size of a list is not fixed like an array.

A **list** declaration consists of the keyword **list**, followed by the data type of the **list** elements between angle brackets, followed by the name(s) of the **list**(s).

```
struct my_s {
  list<int> my_list;
}
```

Example 18—Declaring a list in a struct

Copyright © 2026 Accellera. All rights reserved.

7.9.3.1 List operators

The following operators are defined for **list**s:

Index operator **[]**

Used to access a specific element of a **list**, given an index into the **list**. The index shall be an integral value. See 8.6.2 for more information on the index operator.

Assignment operator =

Creates a copy of the list-type expression on the RHS and assigns it to the **list** on the LHS. See 8.3 for more information on the assignment operator.

Equality operator ==

Evaluates to *true* if the two **list**s are the same size and all elements with corresponding indexes are equal. Two **list**s of different element types are incomparable. See 8.5.3 for more information on the equality operator.

Inequality operator **!=**

Evaluates to *true* if the two **list**s are not the same size or not all elements with corresponding indexes are equal. Two **list**s of different element types are incomparable. See 8.5.3 for more information on the inequality operator.

Set membership operator **in**

The set membership operator can be applied to a **list** to check whether a specific element is currently in the **list**. It evaluates to *true* if the element specified on the left of the operator exists in the **list** collection on the right of the operator. The type of the element shall be the same as the **list**'s element data type. See 8.5.9 for more information on the set membership operator.

foreach *statement*

The **foreach** statement can be applied to a **list** to iterate over the **list** elements within an activity, a constraint or native exec code. See 11.4.3, 13.1.8, and 20.7.8, respectively, for more information on the **foreach** statements in these contexts.

7.9.3.2 List methods

The following methods are defined for **list**s:

pure function int *size()***;**

Returns the number of elements in the **list**.

function void *clear()***;**

Removes all elements from the **list**.

function *data_type delete(***int** *index)***;**

Removes an element at the specified index of type integer and returns the element value. The return value data type is the same as the data type of the **list** elements. If the index is out of bounds, the operation is illegal.

Copyright © 2026 Accellera. All rights reserved.

function void *insert(***int** *index, data_type element)***;**

Adds an element to the **list** at the specified index of type integer. If the index is equal to the size of the **list**, *insert* is equivalent to *push_back()*. If the index is less than the size of the **list**, then elements at and beyond the index are moved by one. If the index is greater than the size of the **list**, the operation is illegal. The inserted element's data type shall be the same as the data type of the **list** elements.

function *data_type pop_front()***;**

Removes the first element of the **list** and returns the element value. This is equivalent to *delete(0)*.

function void *push_front(data_type element)***;**

Inserts an element at the beginning of the **list**. This is equivalent to *insert(0, element)*.

function *data_type pop_back()***;**

Removes the last element of the **list** and returns the element value. This is equivalent to *delete(size()-1)*.

function void *push_back(data_type element)***;**

Appends an element to the end of the **list**. This is equivalent to *insert(size(), element)*.

pure function string *join*(**string** *connector*)*;*

This method only applies to lists of strings. It concatenates the strings from the list using *connector* as the connector between them, and returns the resulting string.

If *connector* is an empty string, *join*() concatenates the strings with no whitespace inserted.

If the list is empty, an empty string is returned. If the list has 1 element, the string value of this element itself is returned and *connector* is ignored.

pure function string *str_from_chars*();

This method only applies to lists of integer types. It returns the string whose characters' (ASCII) codes appear in the list, in the same order. Depending on the specific type of the list elements, each value is implicitly cast to 8 bits that represent the character code.

pure function set<*data_type*> *to_set()***;**

Returns a **set** containing the elements of the **list**. Each element value will appear once. The **set**'s element data type is the same as the data type of the **list** elements. The **set** is unordered.

function void *shuffle();*

Randomly reorders the elements in the **list**.

7.9.3.3 Example

The following example shows use of **list** operators and methods. In this example, an action of type `B` will be traversed six times. There are six elements in `foo_list3`, `foo_list2[0]` is `1` and `4` is in `comp.foo_list1`. Action `A` and action `C` are never traversed.

Copyright © 2026 Accellera. All rights reserved.

```
component pss_top {
  list<bit[16]> foo_list1, foo_list2;

  exec init_up {
    foo_list1 = {1, 2, 3, 4};  // List initialization with aggregate literal
    foo_list2.push_back(1);    // List initialization with push_back
    foo_list2.push_back(4);
  }

  action A{}
  action B{}
  action C{}

  action traverse_list_a {
    list <bit[16]> foo_list3;
    bit[16] deleted;

    exec pre_solve {
      foo_list3 = pss_top.foo_list1; // foo_list3 = {1, 2, 3, 4}
      foo_list3.push_front(0);       // foo_list3 = {0, 1, 2, 3, 4}
      foo_list3.push_back(5);        // foo_list3 = {0, 1, 2, 3, 4, 5}
      foo_list3.insert(0, 1);        // foo_list3 = {1, 0, 1, 2, 3, 4, 5}
      foo_list3[0] = 6;              // foo_list3 = {6, 0, 1, 2, 3, 4, 5}
      deleted = foo_list3.delete(0); // foo_list3 = {0, 1, 2, 3, 4, 5}
    }

    activity {
      if (comp.foo_list1 == comp.foo_list2) // Use of == operator on list
        do A;
      else foreach (e: foo_list3)           // Use of "foreach" on list
        if (comp.foo_list2[0] > 3)          // Use of [] operator on list
          do A;
        else if (4 in comp.foo_list1)       // Use of "in" operator on list
          do B;
        else
          do C;
    }

    exec post_solve {
      foo_list3.clear();                    // foo_list3 = {}
    }
  }
}
```

Example 19—List operators and methods

7.9.3.4 List randomization

When the context containing the **list** attribute is randomized, the elements of the list are randomized. Random-size lists are not supported. Consequently, it is illegal to place a constraint on the **`size()`** method of a list outside an iterative constraint on the same list. The list size is considered to be an invariant inside an iterative constraint. Consequently, the **`size()`** method may be referenced in constraints within an iterative constraint. Example 20 shows declaration of a list with **bit**-type elements and illustrates valid and invalid constraints on the **`size()`** method.

Copyright © 2026 Accellera. All rights reserved.

```
struct S {
   rand list<bit[8]> lst;

   exec pre_solve {              // Initialize the list
      repeat (100) {
         lst.push_back(0);
      }
   }

   constraint {
      lst.size() in [4..100];    // Error: illegal constraint on size()
      foreach (lst[i]) {
         lst[i] == i+lst.size(); // OK: size() is an invariant in foreach
      }
   }
}
```

Example 20—List randomization

7.9.4 Maps

The **map** collection type is used to declare a variable-sized associative array that associates a *key* with an element (or *value*). The keys serve as indexes into the **map** collection. Using a key, an element in the **map** can be assigned or used in an expression. A **map** is unordered.

A **map** is initialized as an empty collection unless the declaration contains an initialization assignment. A **map** that is empty is as if it was assigned an *empty aggregate literal* (**{}**). **Map** elements can be added or removed within **exec** blocks.

A **map** declaration consists of the keyword **map**, followed by the data type of the **map** keys and the data type of **map** elements, between angle brackets, followed by the name(s) of the **map**(s). Both keys and element values may be of any plain-data type. **Maps** are non-randomizable.

```
struct my_s {
  map<int, string> my_map;
}
```

Example 21—Declaring a map in a struct

7.9.4.1 Map operators

The following operators are defined for **map**s:

Index operator **[]**

Used to access a specific element of a **map**, given a key of the specified data type. When used on the LHS in an assignment, the index operator sets the element value associated with the specified key. If the key already exists, the current value associated with the key is replaced with the value of the expression on the RHS. If the key does not exist, then a new key is added to the **map** collection and the value of the expression on the RHS is assigned to the new key's associated **map** entry. Use of a key that does not exist in the **map** to reference an element in the **map** is illegal. See 8.6.2 for more information on the index operator.

Copyright © 2026 Accellera. All rights reserved.

Assignment operator =

Creates a copy of the map-type expression on the RHS and assigns it to the **map** on the LHS. If the same key appears more than once in the expression on the RHS, the last value specified is used. See 8.3 for more information on the assignment operator.

Equality operator ==

Evaluates to *true* if the two **map**s are the same size, have the same set of keys, and all elements with corresponding keys are equal. Two **map**s of different key or element types are incomparable. See 8.5.3 for more information on the equality operator.

Inequality operator **!=**

Evaluates to *true* if the two **map**s are not the same size, do not have the same set of keys, or not all elements with corresponding keys are equal. Two **map**s of different key or element types are incomparable. See 8.5.3 for more information on the inequality operator.

foreach *statement*

The **foreach** statement can be applied to a **map** to iterate over the **map** elements within an activity, a constraint or native exec code. See 11.4.3, 13.1.8, and 20.7.8, respectively, for more information on the **foreach** statements in these contexts.

The set membership operator (**in**) cannot be applied directly to a map. However, it may be applied to the set of keys or the list of values produced by the *keys()* and *values()* methods, respectively, described below.

7.9.4.2 Map methods

The following methods are defined for **map**s:

pure function int *size()*;

Returns the number of elements in the map.

function void *clear()*;

Removes all elements from the map.

function *data_type delete(data_type key)*;

Removes the element associated with the specified key from the **map** and returns the element value. The return value data type is the same as the data type of the **map** elements. The key argument shall have the same type as specified in the **map** declaration. If the specified key does not exist in the **map**, the operation is illegal.

function void *insert(data_type key, data_type value)*;

Adds the specified key/value pair to the **map**. If the key currently exists in the **map**, then the current value is replaced with the new value. The arguments shall have the same types as specified in the **map** declaration.

pure function set<*data_type*> *keys()*;

Returns a **set** containing the **map** keys. The **set**'s element data type is the same as the data type of the map keys. Since each key is unique and no order is defined on the keys, the method returns a **set** collection.

Copyright © 2026 Accellera. All rights reserved.

pure function list<*data_type*> *values()***;**

Returns a **list** containing the **map** element values. The **list**'s element data type is the same as the data type of the map elements. Since element values may not be unique, the method returns a **list** collection. However, the order of the **list** elements is unspecified.

7.9.4.3 Example

The following example shows use of map operators and methods. In this example, an action of type `B` will be traversed four times: `foo_map1` is not equal to `foo_map2`, `foo_map3` has four elements, `foo_map2["a"]` is `1` which is not greater than `3`, and `"b"` exists in `foo_map1`.

Copyright © 2026 Accellera. All rights reserved.

```
component pss_top {
  map<string, bit[16]> foo_map1, foo_map2;
  list<bit[16]>        foo_list1;

  exec init_up {
    foo_map1 = {"a":1,"b":2,"c":3,"d":4};  // Map initialization
                                           // with key/value literal
    foo_map2["a"] = 1;
    foo_map2["b"] = 4;
    foo_list1 = foo_map1.values();
    foreach (foo_map2[i]) foo_list1.push_back(foo_map2[i]);
  }

  action A{}
  action B{}
  action C{}

  action traverse_map_a {
    rand int lower_size;
    map <string, bit[16]> foo_map3;
    set <string> foo_set1;

    exec pre_solve {
      foo_map3 = pss_top.foo_map1; // foo_map3 = {"a":1,"b":2,"c":3,"d":4}
      foo_map3.insert("z",0); // foo_map3 = {"a":1,"b":2,"c":3,"d":4,"z":0}
      foo_map3.insert("d",5); // foo_map3 = {"a":1,"b":2,"c":3,"d":5,"z":0}
      foo_map3.delete("d");   // foo_map3 = {"a":1,"b":2,"c":3,"z":0}
      foo_set1 = foo_map3.keys();
    }
    constraint lower_size < foo_map3.size() + comp.foo_list1.size();
    activity {
      if (comp.foo_map1 == comp.foo_map2)  // Use of == operator on maps
         do A;
      else foreach (foo_map3.values()[i])  // Use of "foreach" on a map
                                           // converted to a list of values
         if (comp.foo_map2["a"] > 3)       // Usage of operator[] on a map
           do A;
         else if ("b" in comp.foo_map1.keys()) // Check whether a key
                                               // is in the map
           do B;
         else
           do C;
    }
  }
}
```

Example 22—Map operators and methods

7.9.5 Sets

The **set** collection type is used to declare a variable-sized unordered set of unique elements of plain-data type. Sets can be created, modified, and queried using the operators and methods described below.

A **set** is initialized as an empty collection unless the declaration contains an initialization assignment. A **set** that is empty is as if it was assigned an *empty aggregate literal* (**{}**). **Set** elements can be added or removed within **exec** blocks; therefore, the size of a set is not fixed like an array.

Copyright © 2026 Accellera. All rights reserved.

A **set** declaration consists of the keyword **set**, followed by the data type of the **set** elements between angle brackets, followed by the name(s) of the **set**(s). **Set**s are non-randomizable.

```
struct my_s {
  set<int> my_set;
}
```

Example 23—Declaring a set in a struct

7.9.5.1 Set operators

The following operators are defined for **set**s:

Assignment operator =

Creates a copy of the set-type expression on the RHS and assigns it to the **set** on the LHS. The same value may appear more than once in the expression on the RHS, but it will appear only once in the **set**. See 8.3 for more information on the assignment operator.

Equality operator ==

Evaluates to *true* if the two **set**s have exactly the same elements. Note that sets are unordered. Two **set**s of different element types are incomparable. See 8.5.3 for more information on the equality operator.

Inequality operator **!=**

Evaluates to *true* if the two **set**s do not have exactly the same elements. Two **set**s of different element types are incomparable. See 8.5.3 for more information on the inequality operator.

Set membership operator **in**

The set membership operator can be applied to a **set** to check whether a specific element is currently within the **set**. It evaluates to *true* if the element specified on the left of the operator exists in the **set** collection on the right of the operator. The type of the element shall be the same as the **set**'s element data type. See 8.5.9 for more information on the set membership operator.

foreach *statement*

The **foreach** statement can be applied to a **set** to iterate over the **set** elements within an activity, a constraint or native exec code. When applied to a set, the **foreach** statement shall specify an *iterator variable* and shall not specify an *index variable*. See 11.4.3, 13.1.8, and 20.7.8, respectively, for more information on the **foreach** statements in these contexts.

7.9.5.2 Set methods

The following methods are defined for **set**s:

pure function int *size()***;**

Returns the number of elements in the **set**.

function void *clear()***;**

Removes all elements from the **set**.

Copyright © 2026 Accellera. All rights reserved.

function void *delete(data_type element)***;**

Removes the specified element from the **set**. The element argument data type shall be the same as the data type of the **set** elements. If the element does not exist in the **set**, the operation is illegal.

function void *insert(data_type element)***;**

Adds the specified element to the **set**. The inserted element's data type shall be the same as the data type of the **set** elements. If the element already exists in the **set**, the method shall have no effect.

function list<*data_type*> *to_list()***;**

Returns a **list** containing the elements of the **set** in an arbitrary order. The **list**'s element data type is the same as the data type of the **set** elements.

7.9.5.3 Example

The following example shows use of **set** operators and methods. In this example, `A` is traversed two times and `B` is traversed three times: `foo_set1` is not equal to `foo_set2`, there are five elements in `foo_set3`, two of the `foo_set3` elements are in `foo_set2`, and `"b"` is in `foo_set1`.

Copyright © 2026 Accellera. All rights reserved.

```
component pss_top {
  set <string> foo_set1, foo_set2;
  list<string> foo_list1;

  exec init_up {
    foo_set1 = {"a","b","c","d"}; // Set initialization with aggregate literal
    foo_set2.insert("a");
    foo_set2.insert("b");
    foo_list1 = foo_set1.to_list();
    foreach (e:foo_set2) foo_list1.push_back(e);
  }

  action A{}
  action B{}
  action C{rand string character;}

  action traverse_set_a {
    rand int lower_size;
    set <string> foo_set3;
    list<string> foo_list2;

    exec pre_solve {
      foo_set3 = pss_top.foo_set1;
      foo_set3.insert("z");
      foo_set3.insert("e");
      foo_set3.delete("d");
      foo_list2 = foo_set3.to_list();
    }

    constraint lower_size < foo_set3.size() + comp.foo_list1.size();

    activity {
      if (comp.foo_set1 == comp.foo_set2)   // Use == operator on sets
        do A;
      else foreach (e:foo_set3)             // Use "foreach" on set
        if (e in comp.foo_set2)             // Use [] operator on set
          do A;
        else if ("b" in comp.foo_set1)      // Use "in" operator on set
          do B;
        else
          replicate (j:foo_list2.size())
            do C with {character == foo_list2[j];};
    }
  }
}
```

Example 24—Set operators and methods

7.10 Reference types

PSS supports a limited form of *reference types* for modeling elements, but does not support references to plain-data types. *References* in PSS are similar in their semantics to class variables in such languages as Java and SystemVerilog. Variables of reference types can be assigned and compared (see more in 8.3 and 8.5.3).

Copyright © 2026 Accellera. All rights reserved.

7.10.1 Syntax

```
reference_type ::= ref entity_type_identifier
entity_type_identifier ::=
    action_type_identifier
  | monitor_type_identifier
  | component_type_identifier
  | flow_object_type
  | resource_object_type
null_ref ::= null
```

Syntax 16—Ref declaration

The following also apply:

a) Reference types, declared with the **ref** modifier, and collections thereof can be used in the declaration of local variables, function parameters, and function return values. In addition, component reference types only (but not action or flow/resource object reference types) and collections thereof can be used to declare fields in the scope of components, actions, and monitors. Reference types shall not be used to declare fields in the scope of flow/resource objects or structs. It shall be illegal to declare static constants of reference types and collections thereof.

b) Fields and instance functions can be accessed through a reference expression in the same way as through an instance path, using the dot (‘.’) operator.

c) An expression of reference type may evaluate to the special value **null**, indicating that it does not reference any entity. It shall be an error to access members of an entity through a null reference. See also 8.3 and 8.5.3.

d) When not initialized, the default value of a reference variable is **null**.

Note that PSS supports special reference fields that are automatically resolved as part of the solving process. They are:

— The context component reference **`comp`** (see 9.1.5)

— Action handles to sub-actions within compound actions (see 11.3.1.1)

— Input and output reference fields of actions (see 9.3)

— Resource claim reference fields (see 9.4)

Copyright © 2026 Accellera. All rights reserved.

7.10.2 Examples

Example 25 demonstrates the use of a reference as a local variable and as a return type of a function. In the body of **action** `call_foo`, a reference to `A` is stored in a local variable, and then used to call function `foo()`. In addition, a reference to `A` is returned from function `choose_A()`, and it is used in turn to call `foo()` on the chosen instance of `A`.

```
component A {
   function void foo();
};

component B {
   A a_arr[5];

   function ref A choose_A(int code) {
         return a_arr[code % 5];
   }

   action call_foo {
      exec body {
         ref A aref = comp.a_arr[3];
         aref.foo();
         comp.choose_A(123).foo();
      }
   };
};
```

Example 25—Use of reference as local variable and function return value

Example 26 shows a component reference field `close_sibling` declared under **component** `my_comp`. In addition, a list of component references field `all_siblings` is declared under `my_comp`. After the construction of the component instance tree, `c1.close_sibling` is equal to **null** because it was not initialized, while `c2.close_sibling` and `c3.close_sibling` contain references to `c1` and `c2`, respectively. The `all_siblings` list of each one of `c1`, `c2`, and `c3` contains the references to the other two `my_comp` instances, respectively. Consequently, the attribute `close_sibling_data` of `c1` is still equal to its default value 0, and `close_sibling_data` of `c2` and `c3` is equal to 10 and 20, respectively, having been assigned in the `init_down` block through the `close_sibling` reference field. The `all_siblings_data` lists of each of `c1`, `c2`, and `c3` contain `{20,30}`, `{10,30}`, and `{10,20}`, respectively. In addition, a component reference field `some_sibling_comp` is declared under **action** `do_it`. In **exec post_solve** it is randomly assigned to the reference to one of this action's comp's `all_siblings` references. This randomly chosen component reference is then used to print a message in **exec body**.

Copyright © 2026 Accellera. All rights reserved.

```
import std_pkg::*;

component my_comp {
   ref my_comp close_sibling;
   int data, close_sibling_data;
   list<ref my_comp> all_siblings;
   list<int> all_siblings_data;

   exec init_down {
      if (close_sibling != null) {
         close_sibling_data = close_sibling.data;
      }
      foreach (sibling: all_siblings) {
         all_siblings_data.push_back(sibling.data);
      }
   }
   action do_it {
      ref my_comp some_sibling_comp;

      exec post_solve {
         int idx;
         randomize idx with {
            idx in [0..comp.all_siblings.size()-1];
         }
         some_sibling_comp = comp.all_siblings[idx];
      }

       exec body {
         message(LOW, "my component's data is %d", comp.data);
         message(HIGH, "my component's randomly chosen sibling data is %d",
                 some_sibling_comp.data);
      }
   }
}

component pss_top {
   my_comp c1, c2, c3;
   exec init_down {
      c1.data = 10;
      c2.data = 20;
      c3.data = 30;
      c2.close_sibling = c1;
      c3.close_sibling = c2;
      c1.all_siblings = {c2,c3};
      c2.all_siblings = {c1,c3};
      c3.all_siblings = {c1,c2};
   }
}
```

Example 26—Use of reference field and null value

Copyright © 2026 Accellera. All rights reserved.

7.11 User-defined data types

The **typedef** statement declares a user-defined type name in terms of an existing data type, as shown in Syntax 17.

7.11.1 Syntax

```
typedef_declaration ::= typedef data_type identifier ;
```

Syntax 17—User-defined type declaration

7.11.2 Example

A **typedef** example is shown in Example 27.

```
typedef bit[32] uint32_t;
```

Example 27—Typedef

7.12 Data type conversion

Expressions of types **int**, **bit**, **bool**, **enum**, or floating-point type can be changed to another type in this list by using a *cast operator*. In addition, an expression of a reference type can be changed to a compatible reference type.

7.12.1 Syntax

Syntax 18 defines the *cast operator*.

```
cast_expression ::= ( casting_type ) expression
casting_type ::=
    integer_type
  | bool_type
  | enum_type
  | float_type
  | reference_type
  | type_identifier
```

Syntax 18—Cast operation

In a *cast_expression*, the *expression* to be cast shall be preceded by the casting data type enclosed in parentheses. The *cast* shall return the value of the *expression* represented as the *casting_type*. A *type_identifier* specified as a *casting_type* shall refer to a numeric, Boolean, enumeration, or reference type.

The following also apply:

a) A numeric, Boolean, or enumeration value can only be cast to another numeric, Boolean or enumeration type. A reference value can only be cast to a compatible reference type.

Copyright © 2026 Accellera. All rights reserved.

b) Any non-zero value *cast* to a **bool** type shall evaluate to *true*. A zero value cast to a **bool** type shall evaluate to *false*. When casting a **bool** type to another type, *false* evaluates to `0` and *true* evaluates to `1`.

c) All numeric expressions (**int** or **bit** with an optional width specifier, **float32** or **float64**) are type compatible, so an explicit *cast* is not required from one to another.

d) When casting from one numeric type to another numeric type, the conversion shall be performed based on the numeric value conversion rules (see 8.7.2).

e) When casting a value to a user-defined **enum** type, the value shall correspond to the result of an implicit cast to the resulting underlying numeric type. When used in a constraint, the domain of a field of **enum** type consists of the values of the **enum** type.

f) A reference value cast to a (direct or indirect) supertype reference or to its own reference type (*upcast*) shall evaluate to the same reference. An explicit cast is not required in this case; an upcast is implicit.

g) A reference value cast to a (direct or indirect) subtype reference (*downcast*) shall evaluate to the same reference if the dynamic value of the reference belongs to the casting type, and shall evaluate to **null** otherwise.

7.12.2 Examples

Example 28 shows the overlap of possible **enum** values (from 7.12.1 (e)) when used in constraints.

```
import std_pkg::*;

enum config_modes_e {UNKNOWN, MODE_A=10, MODE_B=20};
enum foo_e {A=10, B, C};
function bit[32] get_cfg_mode() {return 30;}
        // a new cfg_mode that has not been added to the enum type yet

action my_a {
   config_modes_e top_cfg;
   rand config_modes_e cfg;
   rand foo_e foo;
   constraint cfg == (config_modes_e)11;
                                 // contradiction - no possible value
   constraint cfg == (config_modes_e)foo;
                                 // cfg==MODE_A, the only value in the
                                 // numeric domain of both cfg and foo
   exec pre_solve {
      config_modes_e cfg_mode = (config_modes_e)get_cfg_mode();
      match (cfg_mode) {
         [MODE_A,
          MODE_B] : top_cfg = cfg_mode;
         [UNKNOWN]: print("Unknown configuration mode\n");
         default  : print("Invalid configuration mode = %d\n",
                          (int)cfg_mode);
      }
   }
}
```

Example 28—Overlap of possible enum values

Example 29 shows the casting of `al` from the `align_e` enum type to a 4-bit unsigned value to pass into the `alloc_addr` imported function.

Copyright © 2026 Accellera. All rights reserved.

```
package external_fn_pkg {
   enum align_e {byte_aligned=1, short_aligned=2, word_aligned=4};
   function bit[32] alloc_addr(bit[32] size, bit[4] align);
   buffer mem_seg_s {
      rand bit[32]  size;
      bit[32] addr;
      align_e al;
      exec post_solve {
         addr = alloc_addr(size, (bit[4])al);
      }
   }
}
```

Example 29—Casting of variable to an unsigned type

Example 30 shows reference type casting on the **comp** field of an action.

```
component C {
   action A {}
}

component sub_C: C {
   int a = 17;
}

extend action C::A {
   int b;
   exec post_solve {
      if ((ref sub_C)comp != null) {
         b = ((ref sub_C)comp).a;
      }
   }
}
```

Example 30—Casting of reference type

7.13 Annotation types

Model annotations allow metadata to be attached to elements of a PSS model. Annotations are metadata that provide data about an element in the PSS model but are not part of the model itself. They have no direct effect on the operation of the code they annotate, but they can be used by the compiler or at runtime by PSS tools to influence behavior.

Copyright © 2026 Accellera. All rights reserved.

When applied, an annotation is attached to the next PSS model element. Annotation types capture the attributes of a specific annotation kind. Annotation types may be declared in package scopes.

```
annotation_declaration ::= annotation annotation_identifier [ template_param_decl_list ]
[ annotation_super_spec ] { { annotation_body_item } }
annotation_super_spec ::= : type_identifier
annotation_body_item ::=
  annotation_attr_field
  | compile_assert_stmt
  | annotation_body_compile_if
  | annotation_usage annotation_body_item
  | stmt_terminator
annotation_attr_field ::= [static const] data_declaration
annotation_body_compile_if ::=
  compile if ( constant_expression ) annotation_body_compile_if_item
    [ else annotation_body_compile_if_item ]
annotation_body_compile_if_item ::= { { annotation_body_item } }
```

Syntax 19—Annotation declaration syntax

The following also apply:

a) Annotation attributes may only be of plain data type and collections thereof.
b) Annotation types may only be declared in the package scope
c) Annotations may be extended, much as other types can be.

Annotations are not a datatype and can only be used as specified below.

```
annotation desc_s {
  string desc;
}
```

Example 31—Annotation declaration

Example 31 shows the declaration of an annotation named *desc_s* that contains a string attribute.

An annotation is applied to a PSS model element using the syntax shown in Syntax 20.

```
annotation_usage ::= @ annotation_type_identifier [ annotation_params_list ]
annotation_params_list ::= { annotation_param_item {, annotation_param_item} }
annotation_param_item ::= . identifier = constant_expression
```

Syntax 20—Annotation application syntax

An annotation can be applied to a specific element, or can be a standalone annotation. A standalone annotation is terminated by a single semicolon. An annotation attached to an element is not, itself, terminated with a semicolon. A standalone annotation is attached to a lexical location, while an element

Copyright © 2026 Accellera. All rights reserved.

annotation is attached to the next PSS model element declared in the scope. It is an error if no subsequent element is present in the scope. Annotations, like comments, are advisory. PSS processing tools shall disregard unrecognized annotations.

The following also apply:

a) All initializer expressions shall be constant.

b) Tools shall only perform semantic checking on recognized annotation types.

```
@desc_s {.desc="Configures the IP"}
action config {
  @desc_c {.desc="Scope annotation"};
  // ...
  @desc_c // Error: neither subsequent element in the scope nor a
          // semicolon
}
```

Example 32—Applying annotations

Example 32 shows both element and standalone annotations. The first annotation is an element annotation. The second is a standalone annotation. The final element annotation is considered an error because no PSS element is subsequently declared within the scope nor a semicolon.

Copyright © 2026 Accellera. All rights reserved.

8. Operators and expressions

This section describes the operators and operands available in PSS and how to use them to form expressions.

An *expression* is a construct that can be evaluated to determine a specific value. Expressions may be *primary expressions*, consisting of a single term, or *compound expressions*, combining *operators* with sub-expressions as their *operands.*

The various types of primary expressions are specified in 8.6.

8.1 Syntax

```
expression ::=
    primary
  | unary_operator primary
  | expression binary_operator expression
  | conditional_expression
  | in_expression
unary_operator ::= - | ! | ~ | & | | | ^
binary_operator ::= * | / | % | + | - | << | >> | == | != | < | <= | > | >= | || | && | | | ^ | & | **
assign_op ::= = | += | -= | <<= | >>= | |= | &=
primary ::=
    number
  | aggregate_literal
  | bool_literal
  | string_literal
  | null_ref
  | paren_expr
  | cast_expression
  | ref_path
  | compile_has_expr
paren_expr ::= ( expression )
cast_expression ::= ( casting_type ) expression
```

Syntax 21—Expressions and operators

8.2 Constant expressions

Some constructs require an expression to be a *constant expression*. The operands of a constant expression consist of numeric and string literals, aggregate literals with constant values, named constants (e.g., **static const**, template parameters), bit-selects and part-selects of named constants, enum items, and calls of **pure** functions with constant arguments.

Copyright © 2026 Accellera. All rights reserved.

8.3 Assignment operators

The assignment operators defined by the PSS language are listed in the table below.

Table 9—Assignment operators and data types

Operator token	Operator name	Operand data types
=	Binary assignment operator	Any plain-data type or reference type
+= -=	Binary arithmetic assignment operators	Numeric
&= |=	Binary bitwise assignment operators	Integer
>>= <<=	Binary shift assignment operators	Integer

The *assignment* (**=**) operator is used in the following contexts:

- Attribute initializers for fields and static constants
- Procedural assignment statements
- Procedural variable declarations with an initial value
- Numeric values in enum item declarations
- Default values for function parameters
- Default values for template value parameters
- Struct literal item definitions
- Annotation parameter lists
- Assignments as part of action initializer lists
- Coverage option declarations

The *arithmetic assignment* (**+=**, **-=**), *shift assignment* (**<<=**, **>>=**), and *bitwise assignment* (**|=**, **&=**) operators are used in the context of procedural statements. These compound assignment operators are equivalent to assigning to the left-hand operand the result of applying the leading operator to the left-hand and right-hand operands. For example, **`a <<= b`** is equivalent to **`a = a << b`**.

While these operators may not be used as a part of an expression, they are documented here for consistency.

The type of the right-hand side of an assignment shall be assignment-compatible with the type of the left-hand side. In an aggregate assignment, assignment is performed element by element. In an assignment of a fixed-size array, the left-hand and right-hand sides of the assignment shall have the same size.

In assignment of **struct** types, the right-hand side shall be of the same type as the left-hand side or a derived type thereof. When the left-hand side of an assignment is of **struct** type and the right-hand side is of a type that inherits from the type of the left-hand side, the elements present in the left-hand type are assigned element-by-element while elements only present in the right-hand type are ignored.

In assignment of reference types, the right-hand side shall be one of the following:

- A reference expression of the same type as the left-hand side or a derived type of it
- An instance path to a component of the same type as the left-hand side or a derived type of it
- The value **null**

Copyright © 2026 Accellera. All rights reserved.

Following the assignment of a reference, the left-hand side variable shall point to (be an alias to) the same entity (component, action, monitor, flow/resource object) referred to by the right-hand side (or have the value **null** in case the right-hand side evaluates to **null**).

8.4 Expression operators

The expression operators defined by the PSS language are listed in the table below.

Table 10—Expression operators and data types

Operator token	Operator name	Operand data types	Result data type
`?:`	Conditional operator	Any plain-data type or reference type (condition is Boolean)	Same as operands
`-`	Unary arithmetic negation operator	Numeric	Same as operand
`~`	Unary bitwise negation operator	Integer	Same as operand
`!`	Unary Boolean negation operator	Boolean	Boolean
`& | ^`	Unary bitwise reduction operators	Integer	Bit
`+ - * / **`	Binary arithmetic operators	Numeric	Numeric
`%`	Binary modulus operator	Integer	Integer
`& | ^`	Binary bitwise operators	Integer	Integer
`>> <<`	Binary shift operators	Integer	Integer
`&& ||`	Binary Boolean logical operators	Boolean	Boolean
`< <= > >=`	Binary relational operators	Numeric	Boolean
`== !=`	Binary logical equality operators	Any plain-data type or reference type	Boolean
cast	Data type conversion operator	Numeric, Boolean, enum	Casting type
`in`	Binary set membership operator	Any plain-data type	Boolean
`[`***expression***`]`	Index operator	Array, list, map	Same as element of collection
`[`***expression***`]`	Bit-select operators	Integer	Integer
`[`***expression***`:` ***expression***`]`	Part-select operator	Integer	Integer

8.4.1 Operator precedence and associativity

Operator precedence and associativity are listed in Table 11. The highest precedence is listed first.

Table 11—Operator precedence and associativity

Operator	Associativity	Precedence
`() []`	Left	1 (Highest)

Copyright © 2026 Accellera. All rights reserved.

Table 11—Operator precedence and associativity (Continued)

cast	Right	2
- ! ~ & | ^ (unary)		2
******	Left	3
*** / %**	Left	4
+ - (binary)	Left	5
<< >>	Left	6
< <= > >= in	Left	7
== !=	Left	8
& (binary)	Left	9
^ (binary)	Left	10
| (binary)	Left	11
&&	Left	12
||	Left	13
?: (conditional operator)	Right	14 (Lowest)

Operators shown in the same row in the table shall have the same precedence. Rows are arranged in order of decreasing precedence for the operators. For example, *****, **/**, and **%** all have the same precedence, which is higher than that of the binary **+** and **–** operators.

All operators shall associate left to right with the exception of the conditional (**?:**) and cast operators, which shall associate right to left. Associativity refers to the order in which the operators having the same precedence are evaluated. Thus, in the following example, `B` is added to `A`, and then `C` is subtracted from the result of `A+B`.

```
A + B - C
```

When operators differ in precedence, the operators with higher precedence shall associate first. In the following example, `B` is divided by `C` (division has higher precedence than addition), and then the result is added to `A`.

```
A + B / C
```

Parentheses can be used to change the operator precedence, as shown below.

```
(A + B) / C             // not the same as A + B / C
```

8.4.2 Using aggregate literals in expressions

Aggregate literals (i.e., value list, map, and structure literals, see 4.8) can be used as expression operands. For example, aggregate literals can be used to initialize the contents of aggregate types as part of a variable declaration, in constraint contexts, as foreign language function parameters, and as template-type value parameters. An aggregate literal may not be the target of an assignment.

Copyright © 2026 Accellera. All rights reserved.

When the operands of an assignment or equality operator are a structure aggregate literal and a **struct**-type variable, any elements not specified by the literal are given the default values of the data type of the element. When the operands of an assignment or equality operator are a value list literal and an array, the number of elements in the aggregate literal shall be the same as the number of elements in the array.

In Example 33, a **struct** type is declared that has four integer fields. A non-random instance of that **struct** is created where all field values are explicitly specified. A constraint compares the fields of this **struct** with an aggregate literal in which only the first two struct fields are specified explicitly. Because a **struct** is a fixed-size data structure, the fields that are not explicitly specified in the aggregate literal are given default values–in this case 0. Consequently, the constraint holds.

```
struct s {
  int a, b, c, d;
};
struct t {
  s s1 = {.a=1,.b=2,.c=0,.d=0};
  constraint s1 == {.b=2,.a=1};
}
```

Example 33—Using a structure literal with an equality operator

When an aggregate literal is used in the context of a variable-sized data type, the aggregate literal specifies both size and content.

In Example 34, a **set** variable is compared with an aggregate literal using a constraint. The size of the **set** variable is three, since there are three unique values in the initializing literal, while the size of the aggregate literal in the constraint is two. Consequently, the constraint does not hold.

```
struct t {
  set<int>   s =  {1, 2, 0, 0};
  constraint s == {1, 2}; // False: s has 3 elements, but the literal has 2
}
```

Example 34—Using an aggregate literal with a set

Values in aggregate literals may be non-constant expressions. Example 35 shows use of a **repeat**-loop index variable and a function call in a value list literal.

```
function int get_val(int idx);
import solve function get_val;
struct S {
  list<array<int,2>>  pair_l;

  exec pre_solve {
    repeat(i : 4) {
      array<int,2> pair = {i, get_val(i)};
      pair_l.push_back(pair);
    }
  }
}
```

Example 35—Using non-constant expressions in aggregate literals

Copyright © 2026 Accellera. All rights reserved.

8.4.3 Type inference rules

The *expected type* of an expression shall be inferred according to the rules below. The expected type is used in the resolution of unqualified *enum_item* names (see 7.5) and in the interpretation of aggregate literals (see 8.4.2).

- The type of the expression on the left-hand side of an assignment determines the expected type of the expression on the right-hand side. This includes initialization assignments.
- The type of the formal parameter of a **function** determines the expected type of the respective actual parameter expression (see 20.2). This is true also for **covergroup** instantiations (see 15.2).
- The return type of a **function** determines the expected type of the expression in its **return** statement (see 20.7.5).
- An expression of a known type on the left-hand side of an equality operator (**==**, **!=**) determines the expected type of the right-hand side (see 8.5.3).
- The expected type of a *conditional_expression* (**?:**) determines the expected type of the second and third operands of the expression (see 8.5.8).
- The type of the expression on the left-hand side of a set membership (**in**) operator determines the expected type of the expressions in the *open_range_list*, or the elements of the *collection_expression*, on the right-hand side (see 8.5.9).
- An explicit data type of a **coverpoint** determines the expected type of the coverpoint expression (see 15.3).
- The type (explicit or implicit) of a **coverpoint** determines the expected type of its bin values (see 15.3.3).
- In a *cast_expression*, the specified target type (*casting_type*) determines the expected type of the expression to be cast (see 7.12).

For the purposes of this section, all numeric types are considered to be a single type as all numeric expressions are type compatible (see 7.12). See more on the evaluation of numeric expressions in 8.7.

In Example 36, contextual typing is required to interpret structure literals. Based on the type of the left operand of an equality operator, the structure literal on the right-hand side is interpreted differently in two different constraints within the same **action**.

Copyright © 2026 Accellera. All rights reserved.

```
component my_ip_c {
  struct my_struct { rand int a; };
  action my_op {
    rand my_struct s;
  }
}

component pss_top {
  my_ip_c my_ip;
  struct your_struct { rand int a; };

  action test {
    rand your_struct s;
    constraint s == {.a = 2};    // pss_top::your_struct literal

    my_ip_c::my_op op;
    constraint op.s == {.a = 3}; // my_ip_c::my_struct literal

    activity {
      op;
    }
  }
}
```

Example 36—Contextual typing in structure literal interpretation

Example 37 shows two cases of unqualified *enum item* resolution based on contextual typing—an assignment and a function call. Note that in calling function `print_num()`, whose formal parameter is declared with type **int**, the identifier `ORANGE` cannot be resolved, because the expected type is an **int**. The *enum_item* shall be qualified in this case.

```
enum color_e {RED, GREEN, ORANGE};

function void print_color(color_e c);
function void print_num(int n);

component pss_top {
  enum fruit_e {APPLE, ORANGE};

  exec init_down {
    color_e c = ORANGE;              // OK - expected type is color_e
    print_color(RED);                // OK - same as above
    print_num((int)ORANGE);          // Error - 'ORANGE' unresolved -
                                     // no enum type expected here
    print_num((int)fruit_e::ORANGE); // OK - qualified reference
  }
}
```

Example 37—Contextual typing in enum_item resolution

8.4.4 Operator expression short-circuiting

The logical operators (**`&&`**, **`||`**) and the conditional operator (**`?:`**) shall use *short-circuit evaluation*. In other words, operand expressions that are not required to determine the final value of the operation shall not be

Copyright © 2026 Accellera. All rights reserved.

evaluated. All other operators shall not use short-circuit evaluation. In other words, all of their operand expressions are always evaluated.

8.5 Operator descriptions

The following sections describe each of the operator categories. The legal operand types for each operator are listed in Table 10.

8.5.1 Arithmetic operators

The binary arithmetic operators are given in Table 12.

Table 12—Binary arithmetic operators

a + b	a plus b
a − b	a minus b
a * b	a multiplied by b (or a times b)
a / b	a divided by b
a % b	a modulo b
a ** b	a to the power of b

Integer division shall truncate the fractional part toward zero. The modulus operator (for example, `a % b`) gives the remainder when the first operand is divided by the second, and thus zero when `b` divides `a` exactly. The result of a modulus operation shall take the sign of the first operand. Division or modulus by zero shall be considered illegal.

If either operand of the power operator is of floating-point type, then the result type shall also be of floating-point type. The result of the power operator is unspecified if the first operand is zero and the second operand is negative or if the first operand is negative and the second operand is not an integer value.

Table 13—Power operator rules for integers

	op1 is < −1	**op1 is −1**	**op1 is 0**	**op1 is 1**	**op1 is > 1**
op2 is positive	op1 ** op2	op2 is odd −> −1 op2 is even −> 1	0	1	op1 ** op2
op2 is zero	1	1	1	1	1
op2 is negative	0	op2 is odd −> −1 op2 is even −> 1	undefined	1	0

The unary arithmetic negation operator (**−**) shall take precedence over the binary operators.

8.5.1.1 Arithmetic expressions with unsigned and signed types

bit-type variables are unsigned, while **int**-type variables are signed.

A value assigned to an unsigned variable shall be treated as an *unsigned* value. A value assigned to a signed variable shall be treated as *signed*. Signed values shall use two's-complement representation. Conversions

Copyright © 2026 Accellera. All rights reserved.

between signed and unsigned values shall keep the same bit representation. Only the bit interpretation changes.

8.5.2 Relational operators

Table 14 lists and defines the relational operators. Relational operators may be applied only to numeric operands.

Table 14—Relational operators

a < b	a less than b
a > b	a greater than b
a <= b	a less than or equal to b
a >= b	a greater than or equal to b

An expression using these *relational operators* shall yield the Boolean value *true* if the specified relation holds, or the Boolean value *false* if the specified relation does not hold.

When one or both operands of a relational expression are unsigned, the expression shall be interpreted as a comparison between unsigned values. If the operands are of unequal bit lengths, the smaller operand shall be zero-extended to the size of the larger operand.

When both operands are signed, the expression shall be interpreted as a comparison between signed values. If the operands are of unequal bit lengths, the smaller operand shall be sign-extended to the size of the larger operand.

All the relational operators have the same precedence, and have lower precedence than arithmetic operators.

8.5.3 Equality operators

The *equality operators* rank lower in precedence than the relational operators. Table 15 defines the equality operators.

Table 15—Equality operators

a == b	a equal to b
a != b	a not equal to b

Both equality operators have the same precedence. When the operands are numeric, these operators compare operands bit for bit. As with the relational operators, the result shall be *false* if the comparison fails and *true* if it succeeds.

When one or both operands are unsigned, the expression shall be interpreted as a comparison between unsigned values. If the operands are of unequal bit lengths, the smaller operand shall be zero-extended to the size of the larger operand.

When both operands are signed, the expression shall be interpreted as a comparison between signed values. If the operands are of unequal bit lengths, the smaller operand shall be sign-extended to the size of the larger operand.

Copyright © 2026 Accellera. All rights reserved.

When the operands of an equality operator are of **string** type, both the sizes and the values of the string operands are compared.

Aggregate data (**struct**s and collections) may be compared using equality operators. When the equality operators are applied to aggregate data, both operands shall be of the same type. Aggregate operands are compared element-by-element to assess equality.

The following rules apply to comparison of collections:

- It shall be illegal to compare two fixed-size arrays of different sizes. Variable-sized collections of the same type may be compared, but they shall be considered not equal if they have different sizes.
- Two fixed-size arrays are considered equal if they have the same elements in the same order.
- Two **list**s are considered equal if they have the same size and they have the same elements in the same order.
- Two **map**s are considered equal if they have the same size and the same *key-value* pairs, regardless of order (maps are unordered).
- Two **set**s are considered equal if they have the same size and the same elements, regardless of order (sets are unordered).

The right-hand side of an equality operator may be an aggregate literal of the same type as the left-hand side. The left-hand side of an equality operator may not be an aggregate literal. See more details about collections in 7.9 and about aggregate literals in 4.8 and 8.4.2.

References can be compared with equality operators. The operands may be one of the following:

- Two expressions of the same reference type, or one expression of a reference to a derived type of the other
- One expression of a component reference type, and the other an instance path to a component of the same type, or a derived type of it
- An expression of a reference type and the value **null**

The expression evaluates to *true* if both operands refer to the same entity (component, action, monitor, flow/resource object) or if both evaluate to **null**; otherwise, it evaluates to *false*. Note that these rules apply to variables declared with the **ref** modifier, the built-in **comp** reference, and other reference fields (see 7.10).

8.5.4 Logical operators

The binary operators *logical AND* (**&&**) and *logical OR* (**||**) are logical connective operators and have a Boolean result. The precedence of **&&** is greater than that of **||**, and both have a lower precedence than the relational and equality operators.

The unary *logical negation* operator (**!**) converts a *true* operand to *false* and a *false* operand to *true*.

In procedural contexts, the **&&** and **||** operators shall use short-circuit evaluation as follows:

- The first operand expression shall always be evaluated.
- For **&&**, if the first operand evaluates to *false*, then the second operand shall not be evaluated.
- For **||**, if the first operand evaluates to *true*, then the second operand shall not be evaluated.

Copyright © 2026 Accellera. All rights reserved.

8.5.5 Bitwise operators

The *bitwise operators* perform bitwise manipulations on the operands. Specifically, the binary bitwise operators combine a bit in one operand with the corresponding bit in the other operand to calculate one bit for the result. The following truth tables show the result for each operator and input operands.

Table 16—Bitwise binary AND operator

&	**0**	**1**
0	0	0
1	0	1

Table 17—Bitwise binary OR operator

|	**0**	**1**
0	0	1
1	1	1

Table 18—Bitwise binary XOR operator

^	**0**	**1**
0	0	1
1	1	0

The bitwise unary negation operator (~) negates each bit of a single operand.

Table 19—Bitwise unary negation operator

~	
0	1
1	0

These operators may be applied only to integer operands.

Copyright © 2026 Accellera. All rights reserved.

8.5.6 Reduction operators

The *unary reduction operators* perform bitwise operations on a single operand to produce a single-bit result.

The *unary AND* operator (**&**) returns **1'b1** if all the bits of the operand are **1**, and returns **1'b0** otherwise. The *unary OR* operator (**|**) returns **1'b1** if any bit of the operand is **1**, and returns **1'b0** otherwise. The *unary XOR* operator (**^**) returns **1'b1** if an odd number of bits of the operand are **1**, and returns **1'b0** otherwise.

These operators may be applied only to integer operands.The table below shows the results of applying the three reduction operators to four example bit patterns.

Table 20—Results of unary reduction operations

Operand	&	|	^	Comments
4'b0000	0	0	0	No bits set
4'b1111	1	1	0	All bits set
4'b0110	0	1	0	Even number of bits set
4'b1000	0	1	1	Odd number of bits set

8.5.7 Shift operators

PSS provides two bitwise *shift operators*: shift-left (**<<**) and shift-right (**>>**). The left shift operator shifts the left operand to the left by the number of bit positions given by the right operand. The vacated bit positions shall be filled with zeros. The right shift operator shifts the left operand to the right by the number of bit positions given by the right operand. If the left operand is unsigned or if the left operand has a non-negative value, the vacated bit positions shall be filled with zeros. If the left operand is signed and has a negative value, the vacated bit positions shall be filled with ones. The right operand shall be a non-negative number. These operators may be applied only to integer operands.

8.5.8 Conditional operator

The *conditional operator* (**?:**) is right-associative and is composed of three operands separated by two operators as shown in Syntax 22. The first operand (the *cond_predicate*) shall be of Boolean type. The second and third operands shall be of the same type, and may be of any plain-data or reference type.

```
conditional_expression ::= cond_predicate ? expression : expression
cond_predicate ::= expression
```

Syntax 22—Conditional operator

If *cond_predicate* is *true*, then the operator evaluates to the first *expression* without evaluating the second *expression*. If *false*, then the operator evaluates to the second *expression* without evaluating the first *expression*.

Copyright © 2026 Accellera. All rights reserved.

8.5.9 Set membership operator

PSS supports the *set membership operator* **in**, as applied to value sets and collection data types. Syntax 23 shows the syntax for the set membership operator.

8.5.9.1 Syntax

```
in_expression ::=
    expression in [ open_range_list ]
  | expression in collection_expression [ [ array_slice ] ]
open_range_list ::= open_range_value { , open_range_value }
open_range_value ::= expression [ .. expression ]
collection_expression ::= expression
array_slice ::=
    expression .. expression
  | expression ..
  | .. expression
```

Syntax 23—Set membership operator

The set membership operator returns *true* if the value of the *expression* on the left-hand side of the **in** operator is found in the *open_range_list* or *collection_expression* on the right-hand side of the operator, and *false* otherwise.

The left-hand side shall not be an unqualified *enum_item* (see 7.5) or an aggregate literal (see 4.8). The elements of the right-hand side of the **in** operator shall have a type compatible with the *expression* on the left-hand side. If the left-hand side and the elements of the right-hand side are numeric expressions, they are context-determined as described in 8.7.1.

If the *expression* on the left-hand side is of a scalar type, the right-hand side may be an *open_range_list* or a *collection_expression*. Otherwise, the right-hand side shall be a *collection_expression*.

An *open_range_list* on the right-hand side of the **in** operator shall be a comma-separated list of scalar value expressions or ranges. When specifying a range, the expressions shall be of a numeric or enumeration type. If the left-hand bound of the range is greater than the right-hand bound of the range, the range is considered empty. Values can be repeated; therefore, values and value ranges can overlap. The evaluation order of the expressions and ranges within the *open_range_list* is nondeterministic. If the type is numeric, each range bound expression shall be implicitly cast to the type of the left-hand side.

A *collection_expression* on the right-hand side of the **in** operator shall evaluate to an **array**, **list**, or **set** type that contains elements whose type is compatible with the type of the *expression* on the left-hand side. For example, the *collection_expression* may be a *value_list_literal* or a hierarchical reference to a **set**. The *collection_expression* may also be an array of *modeling elements* (see Clause 9). In this case, the expression on the left-hand side shall be a corresponding **ref** type.

If *array_slice* is used, then the type of *collection_expression* shall be an **array** or **list** (but shall not be a **set**), and it shall not be a *value_list_literal*. The left and/or the right index *expressions* of *array_slice*, if specified, shall be integer expressions whose values are between 0 and the size of the **array**/**list** - 1. In this case, the set membership operator shall return *true* if the value of the left-hand side is found in *collection_expression* at an index that belongs to *array_slice*, as follows:

Copyright © 2026 Accellera. All rights reserved.

— If both the left and the right indexes are specified for *array_slice*, the set membership operator shall return *true* if the value of the left-hand side is found in the **array**/**list** between the given indexes, inclusive.
— If only the left index is specified for *array_slice*, the set membership operator shall return *true* if the value of the left-hand side is found in the **array**/**list** starting from the left index.
— If only the right index is specified for *array_slice*, the set membership operator shall return *true* if the value of the left-hand side is found in the **array**/**list** ending at the right index.
— If the left index is greater than the right index, or the left index is greater than or equal to the **array**/**list** size, or the right index is negative, the slice is empty. In this case, the set membership operator shall return *false*.
— If both the left and right index expressions are specified, and their self-determined types are different, they are context determined and shall be evaluated using the type propagation rules as described in 8.7.

8.5.9.2 Examples

Example 38 constrains the `addr` attribute field to the range `0x0000` to `0xFFFF`.

```
constraint addr_c {
  addr in [0x0000..0xFFFF];
}
```

Example 38—Value range constraint

In the example below, `v` is constrained to be in the combined value set of `values` and the values specified directly in the *open_range_list* `1`, `2`. In other words, the value of `v` will be in `[1,2,3,4,5]`. The variable `values` of type **list** may not be referenced in an *open_range_list*.

```
struct s {
    list<int> values = {3, 4, 5};
    rand int v;
    constraint v in [1,2] || v in values;
}
```

Example 39—Set membership in collection

In the example below, `v` is constrained to be in the range `1`, `2`, and between `a` and `b`. The range `a..b` may overlap with the values `1` and `2`.

```
struct s {
    rand int v, a, b;
    constraint a < b;
    constraint v in [1,2,a..b];
}
```

Example 40—Set membership in variable range

Copyright © 2026 Accellera. All rights reserved.

In the example below, `v` is constrained to be in `arr`, between the index `a` and the index `b`. Note that the constraint `a <= b` would be redundant because every possible solution that satisfies `v` in `arr[a..b]` would implicitly satisfy `a <= b`.

```
struct s {
    rand int v, a, b;
    rand array<int,3> arr;
    constraint {
        //a <= b;  // this condition is redundant and not needed
        a >= 0;
        b <= 2;
    }
    constraint v in arr[a..b];
}
```

Example 41—Set membership in array slice

8.6 Primary expressions

There are several types of primary expressions (or *simple operands*).

The simplest type of primary expression is a reference (simple or hierarchical) to a variable, constant, or template parameter.

In order to select a single bit of an integer variable or integer named constant (e.g., **static const** or template parameter), a *bit-select* shall be used. In order to select a bit range of an integer variable or integer named constant, a *part-select* shall be used.

In order to get a substring of characters within a **string** variable or **string** named constant, a *sub-string operator* shall be used.

A collection variable of plain-data type can be referenced as a primary expression. In order to select an element within a collection, an *index operator* shall be used.

A **struct** variable can be referenced as a primary expression.

A function call is a primary expression.

There are additional types of primary expressions. Formally, an expression is a primary expression if it is a *primary* as defined in B.18 and unparenthesized.

8.6.1 Bit-selects and part-selects

Bit-selects select a particular bit from a named integer variable or constant using the syntax

```
identifier [ expression ]
```

The index may be any integer expression and may be non-constant.

Part-selects select a fixed range of contiguous bits using the syntax

```
identifier [ constant_expression : constant_expression ]
```

Copyright © 2026 Accellera. All rights reserved.

The value of the first *constant_expression* shall be greater than or equal to the value of the second *constant_expression*.

Bit-selects and part-selects may be used as operands of other operators and as targets of assignments. It shall be illegal for a bit-select or a part-select to access an out-of-bounds bit index.

8.6.2 Selecting an element from a collection (indexing)

The *index operator* (`[]`) is applied to an **array**, **list**, or **map** collection to select a single element. In the case of an **array** or a **list**, the index shall be an integer expression whose value is between 0 and the size of the **array**/**list** - 1. In the case of a **map**, the index shall be of the same type as that of the key in the **map** declaration.

An indexed collection may be used as an operand of other operators and as a target of assignments.

In the case of an **array** or a **list**, it shall be illegal to access an out-of-bounds index. In the case of a **map**, it shall be illegal to read an element whose key does not appear in the **map**. An assignment to a **map** element whose key does not currently appear in the **map** shall add that key and value pair to the **map**.

8.6.3 The sub-string operator

The *sub-string operator* is applied to a **string** and returns a string that is equal to the sub-string, starting and ending at the specified positions, using the syntax:

```
identifier [ string_slice ]
```

The syntax for the `string_slice` operator is specified in Syntax 24.

```
string_slice ::=
    expression [ .. expression ]
  | expression ..
  | .. expression
```

Syntax 24—String slice operator

All *expressions* are of integer types. The *expression* preceding denotes the starting position, and the *expression* following denotes the ending position. If "**..**" is used but the starting or the ending position is not specified, the sub-string starts from the first character in the string or ends at the last character in the string, respectively. If "**..**" is not used and only one *expression* is specified, it denotes both the starting and the ending position, i.e., the sub-string contains one character.

The starting and ending positions shall be between 0 and the size of the string less than 1, and the starting position shall not be greater than the ending position.

8.7 Evaluation rules for numeric expressions

The type of a numeric expression is determined by the operands involved in the expression and the context in which the expression appears. These rules shall be in place in any context; in particular, they shall apply to expressions used in constraints, activity statements, procedural **exec** blocks or functions, and so on. In the case of procedural code, they shall apply on both the solve platform and the target platform.

Copyright © 2026 Accellera. All rights reserved.

8.7.1 Rules for expression types

A *self-determined expression* is one whose type is solely determined by the expression itself.

A *context-determined expression* is one whose type is determined both by the expression itself and by its usage context. For example, if an expression is used as an operand in an arithmetic addition expression, its type depends both on its own operands and on the second operand of the addition expression, i.e.; the type of the right-hand expression of an assignment depends on itself and the type of the left-hand side (see 8.7.2).

Table 21 shows how the type of primary expressions (see 8.6) is determined.

Table 22 shows how the type of compound (non-primary) expressions is determined, and the type propagation rules for the expression operands. In general, where the type of an expression is propagated to an operand, that operand is context-determined.

Both Table 21 and Table 22 determine the self-determined type of an expression. When this expression itself serves as a context-determined operand in another expression, its final type is determined according to Table 22. When this expression is a source expression in an assignment-like context (see 8.7.2), its final type is determined according to the description in 8.7.2.

The actual numeric value conversions as part of evaluating an expression shall be performed as follows:

- Converting a context-determined simple operand value from one integer type to another (propagated) integer type: if the propagated type is signed, the value shall be sign-extended to the propagated bit size; if the propagated type is unsigned, the value shall be zero-extended to the propagated bit size.
- Converting an integer operand value to a floating-point type: the integer value shall be converted to the floating-point representation of the same integer value.

Table 21—Types of primary expressions

Expression	Type
Unsized decimal or octal constant number	**int[N]**, N is the minimal needed number of bits, but at least 32
Unsized hexadecimal or binary constant number	**bit[N]**, N is the minimal needed number of bits, but at least 32
Sized constant number	As specified
Floating-point constant number	**float64**
Variable or hierarchical path	As declared
Function call	The function return type
Bit select	**bit**
Part select	Bit size as per range, unsigned
Cast operator	The specified type

Copyright © 2026 Accellera. All rights reserved.

Table 22—Types of compound expressions and their propagation to operands

Operators	Expression type	Type propagation

Binary arithmetic operators:
`+ - * / %`
Trinary conditional operator (with two numeric operands):
`? :`
	— At least one floating-point operand:
	—float64
	— Two integer operands:
	—Bit size: larger of the two.
	—Signed if both are signed; otherwise, unsigned.
		— One floating-point and one integer operand:
		—The integer operand shall be treated as self-determined, and then converted to floating-point just before the operator is applied.
		— Two integer operands:
		—Propagated to both operands.

Binary power operator: `**`	— At least one floating-point operand: **—float64** — Two integer operands: —Same as the left-hand size operand.	Propagated to left-hand side unless floating-point. Right-hand side is self-determined.
Binary bitwise operators: `& | ^`	Bit size: larger of the two. Signed if both are signed; otherwise, unsigned.	Propagated to both operands.
Binary shift operators: `>> <<`	Same as the left-hand side operand.	Propagated to left-hand side. Right-hand side is self-determined.
Binary equality and relational operators: `== != < <= > >=`	**bool**	Same rules as for binary arithmetic operators: — Two integer operands. Propagated type determined as follows: —Bit size - larger of the two. —Signed if both operands are signed; otherwise, unsigned. — One floating-point and one integer operand: —The integer operand shall be treated as self-determined and then converted to floating-point just before the operator is applied.
Binary and unary logical operators: `&& || !`	**bool**	All operands are self-determined.
Unary arithmetic and bitwise operators: `+ - ~`	Same as the operand.	Propagated to the operand.
Unary reduction operators: `& | ^`	**bit**	Operand is self-determined.

Copyright © 2026 Accellera. All rights reserved.

Table 22—Types of compound expressions and their propagation to operands (Continued)

Operators	Expression type	Type propagation
Set membership operator: **in** (right-hand side is a collection of a known type)	**bool**	— Both left-hand side operand and right-hand side item type are integers. Propagated type to both is determined as follows: —Bit size: larger of the two. —Signed if both types are signed; otherwise, unsigned. — Left-hand side is floating-point. Right-hand side item type is an integer: —Right-hand side collection item values shall be converted to floating-point. — Left-hand side is an integer. Right-hand side item type is floating-point: —The left-hand side operand shall be treated as self-determined and then converted to floating point just before the operator is applied.
Set membership operator: **in** (right-hand side is a value list literal or a range list)	**bool**	Considering multiple items as separate expressions: the left-hand side operand and every item on the right-hand side. Some of them can be integers, some floating point. — If some items are integers, propagated integer type is determined as follows: —Bit size - the largest among all. —Signed if all are signed; otherwise, unsigned. — If some items are floating-point, the integer items shall be converted to floating point just before the operator is applied (but after being evaluated according to the propagated type).

8.7.2 Assignment-like contexts

An *assignment-like* context is a context in which a *source expression* of a numeric type shall be explicitly converted to a specific other numeric type, referred to as the *target type*. The assignment-like contexts are listed in Table 23.

In assignment-like contexts, the type of the source expression is determined as follows:

— If the source expression is floating-point, its type is self-determined.

Copyright © 2026 Accellera. All rights reserved.

— If the target type is a floating-point type (**float32** or **float64**), the source expression type is self-determined.

— Otherwise (both the target type and the source expression type are integer types), if the bit size of the target type is larger than the bit size of the source expression, this bit size is propagated to the source expression. Signedness of the source expression type is self-determined and does not depend on the target type.

The actual conversion of the source expression's value into the target type shall be performed as follows:

— If both types are integers, the source expression's value shall be truncated or extended as necessary to match the bit size of the target type. When the bit size of the target type is larger than the bit size of the source expression's type, sign-extension shall be used if the source expression's type is signed, or zero-extension if the source expression's type is unsigned.

— If the source expression's type is an integer, and the target type is floating-point, the integer value of the expression shall be converted to the floating-point representation of the same integer value.

— If the source expression's type is floating-point, and the target type is an integer, the floating-point value shall be treated as a signed number, and the fractional part of the number shall be truncated.

Table 23—Assignment-like contexts

Context	Source expression	Target type	Comments
Assignment operator (8.3)	Right-hand side	Left-hand side type	Assignment operator is used in the context of different syntactic constructs (see details in 8.3), including some in which the right-hand side is always a constant expression. However, the type conversion rules equally apply to all of them.
Cast operator (7.12)	The operand *expression*	The *casting_type*	
Function call parameters (20.2.7)	The actual parameter	The declared parameter type	Does not apply to generic parameters: for such parameters, the actual parameter expression is self-determined.
return (20.7.5), in a non-void function	The operand *expression*	The function return type	
Value list literal (4.8.2)	An element in the literal	The element type of the context collection type	— Not applicable when there is no context type. — Not applicable when used as the right-hand side of a set membership operator.
Map literal (4.8.3)	A key/value in the literal	The key/value type of the context map type	Not applicable when there is no context type.

Copyright © 2026 Accellera. All rights reserved.

Table 23—Assignment-like contexts (Continued)

Context	Source expression	Target type	Comments
Value parameters in template instantiations (10.4)	The actual parameter	The declared parameter type	
coverpoint declaration (15.3)	The coverage point expression	The declared coverage point type	Applies only when a type is specified. Otherwise, the expression is self-determined.

8.7.3 Examples

Example 42 illustrates how the type propagation rules work when operands of different bit sizes are involved.

```
import std_pkg::*;
component pss_top {
  exec init_up {
    bit[8] x = 0xf0;
    bit[12] y = 0xff0;

    print("x+y=%x\n", x+y); // the value of x is cast to 12 bits
                           // printed result is 0xe0 - after overflow

    bit[16] z = x+y; // result is 0x10e0
                     // the values of x and y are cast to 16 bits
    print("z=%x\n", z);
  }
}
```

Example 42—Bit size propagation

Example 43 illustrates how the type propagation rules work when operands of different signedness are involved. The example shall print "`x is smaller than y`" because the negative value of `y` is converted to unsigned, which is larger than the value of `x`.

```
import std_pkg::*;
component pss_top {
  exec init_up {
    bit[32] x = 1;
    int y = -10;
    if (x < y)
      print("x is smaller than y\n");
    else
      print("x is not smaller than y\n");
  }
}
```

Example 43—Signedness propagation

Copyright © 2026 Accellera. All rights reserved.

Example 44 illustrates how the type propagation rules work when both integer and floating-point operands are involved. In the expression `x + y + z`, the largest bit size 32 is propagated to all operands. Therefore, `x + y` evaluates to `5`, and the final result is `15`.

In the expression `x + y + f`, the sub-expression `x + y` is treated as self-determined, and according to bit size `2`, it evaluates to `1`. Then, the value `1` is converted to floating-point, and the final result of the expression is `11`.

```
import std_pkg::*;
component pss_top {
  exec init_up {
    bit[2] x = 3;
    bit[2] y = 2;
    bit[32] z = 10;
    float64 f = 10;

    print("x + y + z = %d\n", x + y + z); // prints 15

    print("x + y + f = %f\n", x + y + f); // prints 11
  }
}
```

Example 44—Expressions with integer and floating-point operands

Example 45 illustrates how the type propagation rules are applied to the set membership operator **in**.

In the first use of **in**, the left-hand side expression `x + y` is treated as `bit[32]` because the bit size of the list elements is 32. Therefore, the value of the expression is 5, and the condition holds.

In the second use of **in**, all three addition expressions are treated as `bit[3]` as per the type propagation rules where `bit[2]` and `int[3]` are involved. Therefore:

- `x+y` is evaluated to 5
- `a+b` is evaluated to 4
- `c+d` is evaluated to 6: value -4 of `c`, after being converted from **`int[3]`** to **`bit[3]`**, becomes 4

and the condition holds because it is then treated as: `3'd5 in [3'd4..3'd6]`.

Copyright © 2026 Accellera. All rights reserved.

```
import std_pkg::*;
component pss_top {
    exec init_up {
        bit[2] x=3, y=2;
        list<int> my_list = {4,5,6};
        if ((x + y) in my_list)     // the condition holds
            print("(x + y) is in the list\n");
        else
            print("(x + y) is not in the list\n");

        int[3] a=1, b=3, c=-4, d=2;
        if ((x + y) in [a+b..c+d])  // the condition holds
            print("(x + y) is in the range\n");
        else
            print("(x + y) is not in the range\n");
    }
}
```

Example 45—Type propagation with set membership operator

Copyright © 2026 Accellera. All rights reserved.

9. Modeling element types

PSS enables the creation of portable, reusable verification scenarios across simulation, emulation, and hardware platforms. At the heart of PSS are its modeling elements—the fundamental constructs used to describe both the structure and behavior of stimulus models:

- **Components**: Structural containers that define hierarchy and encapsulate behavior.
- **Actions**: Behavioral units that describe operations, transactions, and control logic.
- **Flow Objects**: Represent data/control flow and define preconditions and postconditions for actions.
- **Resource Objects**: Shared or exclusive entities used to coordinate access and synchronization.
- **Monitors**: Represent the scenario to be observed during coverage (see 16.1).

Each of these elements plays a distinct role in expressing verification intent and system behavior. These are described in detail in the following sections.

In this document, the term *modeling elements* is used to refer collectively to **action or monitor handles**, **component instances**, and **references to flow and resource objects**.

All modeling elements are represented by a non-plain data type. Each object instance in the PSS model is assigned by the PSS tool a unique identifier accessible via the built-in attribute *uid*. This identifier ensures consistent referencing and traceability throughout the model. The following rules apply:

- The type of the *uid* field is bit[32].
- The *uid* is globally unique across all modeling element types within the model; no two instances may share the same identifier.

9.1 Components

Components serve as a mechanism to encapsulate and reuse elements of functionality in a portable stimulus model. Typically, a model is broken down into parts that correspond to roles played by different actors during test execution. Components often align with certain structural elements of the system and execution environment, such as hardware engines, software packages, or testbench agents.

Components are structural entities, defined per type and instantiated under other components (see Syntax 25). Component instances constitute a hierarchy (tree structure), beginning with the top or root component, called `pss_top` by default, which is implicitly instantiated. Each component instance has a unique hierarchical path name, and may also contain data attributes, but not constraints. Components may also encapsulate function declarations (see 20.2.1) and imported class instances (see 20.4.3). In addition, components may be derived from other components via inheritance, or a component may be extended to add elements to the component type (see Clause 17).

Copyright © 2026 Accellera. All rights reserved.

9.1.1 Syntax

```
component_declaration ::= [ pure ] component component_identifier [ template_param_decl_list ]
   [ component_super_spec ] { { component_body_item } }
component_super_spec ::= : type_identifier
component_body_item ::=
     override_declaration
   | component_data_declaration
   | component_pool_declaration
   | action_declaration
   | override_action_declaration
   | abstract_action_declaration
   | generic_constraint_declaration
   | object_bind_stmt
   | exec_block
   | struct_declaration
   | enum_declaration
   | covergroup_declaration
   | function_decl
   | import_class_decl
   | procedural_function
   | import_function
   | export_function
   | target_template_function
   | export_action
   | typedef_declaration
   | import_stmt
   | extend_stmt
   | compile_assert_stmt
   | attr_group
   | component_body_compile_if
   | annotation_usage component_body_item
   | monitor_declaration
   | cover_stmt
   | stmt_terminator
```

Syntax 25—Component declaration

9.1.2 Example

For an example of how to declare a component, see Example 46.

Copyright © 2026 Accellera. All rights reserved.

```
component uart_c { ... };
```

Example 46—Component

9.1.3 Components as namespaces

Component types serve as namespaces for their nested types, e.g., **action**, **monitor**, and **struct** types defined under them. The fully-qualified name of nested types is of the form '*package-namespace*::*component-type*::*nested-type*'. References to nested types in a component shall follow the name resolution rules defined in 18.3.

For an example of how to use a component as a namespace, see Example 47.

```
component usb_c {
  action write {...}
}
component uart_c {
  action write {...}
}
component pss_top {
  uart_c s1;
  usb_c s2;
  action entry {
    uart_c::write wr; //refers to the write action in uart_c
    ...
  }
}
```

Example 47—Namespace

In Example 48 below, a **component** `C1` is declared in a **package**. That **component** is instantiated in **component** **`pss_top`**, and an **action** within **component** `C1` is traversed in **action** `pss_top::entry`. In the traversal of **action** `P::C1::A`, the qualified name elements are the following:

— *package-namespace*: `P`
— *component-type*: `C1`
— *class-type*: `A`

Copyright © 2026 Accellera. All rights reserved.

```
package P {
    component C1 {
        action A {}
    }
}

component pss_top {
    P::C1 c1;

    action entry {
        activity {
            do P::C1::A;
        }
    }
}
```

Example 48—Component declared in package

9.1.4 Component instantiation

Components are instantiated under other components as their fields, much like data fields of structs, and may be arrays, possibly multidimensional, thereof.

9.1.4.1 Semantics

a) Component fields are non-random; therefore, the **rand** modifier shall not be used. Component data fields represent configuration data that is accessed by actions and monitors declared in the component and by cover statements invoked in this component. To avoid infinite component instantiation recursion, a component type and all template specializations thereof shall not be instantiated under its own sub-tree.

b) In any model, the component instance tree has a predefined root component, called **`pss_top`** by default, but this may be user-defined. There can only be one root component in any valid scenario.

c) Other components, actions, or monitors are instantiated (directly or indirectly) under the root component. See also Example 49.

d) Plain-data fields may be initialized using a constant expression in their declaration. Data fields may also be initialized via an **exec init_down** or **init_up** block (see 20.1.2), which overrides the value set by an initialization assignment. The component tree is elaborated to instantiate each component and then the **exec init_down** and **init_up** blocks are evaluated hierarchically. See also Example 280 and Example 281 in 20.1.3.

e) Component data fields are considered immutable once construction of the component tree is complete. Actions can read the value of these fields, but cannot modify their value. Fields declared as mutable are an exception and may be modified after the construction of the component tree (see 9.1.7). Component data fields are accessed from actions relative to the **`comp`** field, which is a handle to the component context in which the action is executing. See also Example 282 (and 20.1).

f) It shall be illegal to access static component members using the **`comp`** handle.

g) It shall be illegal to reference non-static context component members from **struct** types declared within the component.

h) Any non-static component member may be referred to with a full hierarchical path starting with the root component.

Copyright © 2026 Accellera. All rights reserved.

9.1.4.2 Examples

Example 49 depicts a component tree definition. In total, there are two instances of `multimedia_ss_c` (instantiated as an array in **`pss_top`**), twelve instances of `codec_c` (six from each two-dimensional array declared in `multimedia_ss_c`), and 24 instances of `vid_pipe_c` (two in each element of the `codec_c` array).

```
component vid_pipe_c { ... };

component codec_c {
   vid_pipe_c pipeA, pipeB;
   action decode { ... };
};

component multimedia_ss_c {
   codec_c codecs[34][2];
};

component pss_top {
   multimedia_ss_c multimedia_ss[2];
};
```

Example 49—Component instantiation

Example 50 shows some legal and illegal accesses to component functions and attributes.

Copyright © 2026 Accellera. All rights reserved.

```
component my_comp_c {
  int f;
  struct S {
    rand int g;
    exec post_solve {
      g = f; // ILLEGAL: S may not refer to instance-specific field of
             // my_comp_c. NOTE: 'g = my_comp_c::f;' would be legal if
             // f were 'static const int f'
    }
  }
  solve function void print_f() {
    print ("%d", f);
  }
  action A_a {
    rand S s;
    exec post_solve {
      comp.print_f(); // comp handle required to access 'print_f'
    }
  }
}

component pss_top {
  my_comp_c comp1, comp2, comp3;
  exec init_up {
    comp1.f = 6; comp2.f = 7; comp3.f = 8;
  }
  action entry_a {
    activity {
      do my_comp_c::A_a;
    }
  }
}
```

Example 50—Component attribute and function access

9.1.5 Component references

Component references can be used in several locations. Each action instance is associated with a specific component instance of its containing component type, the component-type scope where the action is defined. The component instance is the “actor” or “agent” that performs the action. Only actions defined in the scope of instantiated components can legally participate in a scenario. Attributes and collections of component-reference type may also be declared in components and actions. An action may also contain user-defined component reference fields or collections thereof (see 7.10.1 and Example 26).

The component instance with which an **action** is associated is referenced via the built-in field **`comp`**. The value of the **`comp`** field can be used for comparisons of references (see 8.5.3). Unlike user-defined reference variables, the **`comp`** field is assigned automatically as part of the solving process (see 13.4.5) and may not be assigned by the user. The static type of the **`comp`** field is the **ref** type of the action’s context **component**. Consequently, attributes and sub-components of the containing component may be referenced via the **`comp`** field using relative paths.

9.1.5.1 Semantics

A compound action can only instantiate sub-actions that are defined in its containing component or defined in component types that are instantiated in its containing component's instance sub-tree. In other words,

Copyright © 2026 Accellera. All rights reserved.

compound actions cannot instantiate actions that are defined in components outside their context component hierarchy. Similarly, monitors and cover statements cannot reference actions or other monitors that are defined in components outside their context component hierarchy. This maximizes the reusability of components in other contexts.

9.1.5.2 Examples

Example 51 illustrates the need to define a sub-action in a containing component or its sub-tree. In action `graphics::gr_a`, the traversal of `bus_c::write` is illegal since the component `bus_c` is not instantiated in the action's containing component (`graphics`).

```
component bus_c {
  import bar_pkg::*;
  action write{input bar_s b;...} // bar_s is a stream
}

component graphics {
  import bar_pkg::*;
  action foo {output bar_s b;...}
  action gr_a {
    activity {
      parallel {
        do bus_c::write;  // illegal
        do foo;
      }
    }
  }
}

component pss_top {
  import bar_pkg::*;
  bus_c a0;
  graphics g;
  pool bar_s bar_p;
  bind bar_p *;
}
```

Example 51—Illegal traversal of an action outside of the containing component hierarchy

Example 52 demonstrates the use of the **`comp`** reference. The constraint within the `decode` action forces the value of the action's `mode` bit to be `0` for the `codecs[0]` instance, while the value of `mode` is randomly selected for the other instances. The sub-action type `program` is available on both sub-component instances, `pipeA` and `pipeB`, but in this case is assigned specifically to `pipeA` using the **`comp`** reference.

See also 13.1.4.

Copyright © 2026 Accellera. All rights reserved.

```
component vid_pipe_c { ... };
component codec_c {
  vid_pipe_c pipeA, pipeB;
  bit mode1_enable;
  action decode {
    rand bit mode;
    constraint set_mode {
      comp.mode1_enable==0 -> mode == 0;
    }
    activity {
      do vid_pipe_c::program with { comp == this.comp.pipeA; };
    }
  };
};
component multimedia_ss_c {
  codec_c codecs[2];
  exec init_up {
    codecs[0].mode1_enable = 0;
    codecs[1].mode1_enable = 1;
  }
};
```

Example 52—Using the comp attribute in constraints

9.1.5.3 Instance component reference

```
component_data_declaration ::=
  [ access_modifier ] [ component_data_decl_qualifier ] data_declaration
component_data_decl_qualifier ::=
    static const
  | mutable
  | instance
```

Syntax 26—Component ref instance qualifier

A non-null instance-qualified component reference field has the same capabilities as a non-reference component field. The instance qualifier may be applied to component fields whose type is component ref or a collection thereof.

Component-ref attributes in a component can be used in multiple ways—some that may result in reference cycles within the component tree. Reference cycles are permitted when a component ref is used by procedural code. In contrast, reference cycles shall not exist in collections of component fields used as an action-traversal context.

Instance component reference fields have the same capabilities as non-reference component fields when assigned a non-null value. These fields are marked with the instance qualifier, which designates their contextual role. Specifically:

- They can serve as **action-traversal context** within the component hierarchy.
- Static pool-binding directives do not apply to instance-qualified component reference fields.
- Instance component references of address-space and executor types have the same capabilities as instance fields of the respective types.

Copyright © 2026 Accellera. All rights reserved.

The following rules apply to instance component reference fields:

a) The instance qualifier may only be applied to component fields whose type is a component reference or a list or array of component references.

b) Reference cycles shall not exist in the component tree when instance reference fields are used as they can compromise traversal logic and system integrity.

9.1.5.4 Instance-qualified component ref examples

```
component C {
  action A { }
}
component D : C{
  action B { }
}
component F {
  C c1;
  ref E e1;
}
component E {
  instance ref D  d1;
  instance ref C  c1;
  instance ref F f1;
  instance list<ref D> d_l;
           ref C  c2;
  action E_a {
    activity {
      do C::A;
      do D::B;
    }
  }
}
```

Example 53—Use of instance-qualified component ref fields

Example 53 shows component ref fields that are instance-qualified and unqualified. When traversing action `C::A`, the following apply:

— `d1` and the component references in `d_l` are valid action-traversal contexts, provided the refs are not null.

— `c1` is a valid action-traversal context for `C::A`, provided it is not null.

— `c1` is a valid action-traversal context for `D::B`, provided it is not null and is assigned to an instance of `component D`

— `c2` is never a valid action-traversal context because it is a non-instance-qualified reference.

Copyright © 2026 Accellera. All rights reserved.

```
buffer my_buffer { }
component ab_c {
  action actA_a {
    output my_buffer buf_out;
  }
}
component delegator_c {
  pool my_buffer my_buffer_p;
  bind my_buffer_p *;
  instance ref ab_c ab0_r;
  ab_c ab1;
}
component pss_top {
  pool my_buffer my_buffer_p;
  bind my_buffer_p ab0.*;
  ab_c ab0;
  delegator_c   del;
  init_down {
    del.ab0_r = ab0;
  }
}
```

Example 54—Static bind and component instance-ref

In Example 54 above, a static pool bind in `pss_top` is applied to sub-component instance `ab0`. A static pool bind in `delegator_c` is applied to all sub-scopes. Because static binds have no effect on instance component refs, the static bind in `delegator_c` does not apply to the instance component ref `ab0_r`. Consequently:

— Action `ab_c::actA_a.buf_out` is bound to `pss_top.my_buffer_p` when associated with `pss_top.ab0` or `pss_top.del.ab0_r`.

— Action `ab_c::actA_a.buf_out` is bound to `pss_top.del.my_buffer_p` when associated with `pss_top.del.ab1`.

9.1.6 Mutable component attributes

```
component_data_declaration ::=
   [ access_modifier ] [ component_data_decl_qualifier ] data_declaration
component_data_decl_qualifier ::=
    static const
  | mutable
  | instance
```

Syntax 27—Mutable qualifier

A component is considered immutable after its **init_up exec** block has been evaluated as it represents a structural element of the system. To support data that is not inherently part of the system, mutable attributes are available.

Component fields declared with the *mutable* qualifier represent data that is not part of the component's immutable structure and may be modified during solve-time execution, including during the execution of solve-time **exec** blocks or as part of activity solving. Such fields shall not be modified directly or indirectly

Copyright © 2026 Accellera. All rights reserved.

from target exec blocks. The PSS tool shall guarantee that the execution of solve-time exec blocks is atomic to ensure there are no race conditions and the data is updated and read safely.

The following also apply:

a) For aggregate data types, the *mutable* qualifier propagates to all contained elements. For example, if a field of a struct type is declared mutable, all fields within that struct are considered mutable.

b) The *mutable* qualifier shall not be applied to component fields or instance reference fields. The mutability of fields within a referenced or instantiated component is determined solely by their own *mutable* qualification.

c) When applied to a non instance reference field, the *mutable* qualifier applies to the reference itself, allowing the field to be updated to refer to a different component instance. The fields of the referenced component may be modified only if those fields are themselves declared mutable.

Example 55 illustrates the use of mutable attributes to accumulate state across repeated executions of an action. Each execution of `my_comp_c::A_a` produces a value `s` and contributes it to the parent `pss_top` instance by invoking `add_to_sum`. It also records the corresponding component by adding it to `used_comps`. Because `total_sum` and `used_comps` are declared as mutable, they are updated incrementally as the `repeat` statement executes the action multiple times. In contrast, the name field of each `my_comp_c` instance is assigned in `init_down` and remains unchanged thereafter as it is not declared mutable.

Copyright © 2026 Accellera. All rights reserved.

```
component my_comp_c {
  string name;
  action A_a {
    rand int in [1..100] s;
    ref pss_top top_comp;
    exec post_solve {
      // update sum if needed
      if (top_comp != null) {
        top_comp.add_to_sum(s);
        if (!(comp in top_comp.used_comps))  {
          top_comp.used_comps.push_back(comp);
        }
      }
    }
  }
}

component pss_top {
  my_comp_c comp_arr[3];
  mutable int total_sum;
  mutable list<ref my_comp_c> used_comps;

  exec init_down {
    foreach (c: comp_arr [i]) {
      //name cannot be changed in later phase because it's immutable
      c.name = std_pkg::format("comp_arr[%d]", i);
    }
  }

  solve function void add_to_sum(int sum) {
    total_sum = total_sum + sum ;
  }

  action test {
    rand int in [1..10] rep;
    activity {
      repeat (rep)
        do my_comp_c::A_a {.top_comp = comp;};
    }

    exec pre_body {
      std_pkg::print("Total sum: %d\n", comp.total_sum);
      foreach (c: comp.used_comps) {
        std_pkg::print("comp.name: %s\n", c.name);
      }
    }
  }
}
```

Example 55—Mutable qualifier

9.1.7 Pure components

Pure components are restricted types of components that provide PSS implementations with opportunities for significant optimization of storage and initialization. Pure components are used to encapsulate realization-level functionality and cannot contain scenario model features. *Register structures* are one possible application for pure components (see 21.14).

Copyright © 2026 Accellera. All rights reserved.

The following rules apply to pure components, that is, component types declared with the **pure** modifier:

a) In the scope of a **pure component**, it shall be an error to declare **action** and **monitor** types, **pool** instances, **pool** binding directives, non-static data attributes, instances of non-pure **component** types, **exec** blocks, or to specify **cover** statements.

b) A **pure component** may be instantiated under a non-pure **component**. However, a non-pure **component** may not be instantiated under a **pure component**.

c) A **pure component** may not be derived from a non-pure **component**. However, both a **pure component** and a non-pure **component** may be derived from a **pure component**.

An example of the use of pure components is shown in Example 56.

```
pure component my_register {
   function bit[32] read();
   function void write(bit[32] val);
};

pure component my_register_group {
   my_register regs[10];
};

component my_ip {
   my_register_group reg_groups[100]; // sparsely-used large structure
};
```

Example 56—Pure components

9.2 Actions

Actions are a key abstraction unit in PSS. Actions serve to decompose scenarios into elements whose definitions can be reused in many different contexts. Along with their intrinsic properties, actions also encapsulate the rules for their interaction with other actions and the ways to combine them in legal scenarios. Atomic actions may be composed into higher-level actions, and, ultimately, to top-level test actions, using activities (see Clause 11). The *activity* of a compound action specifies the intended schedule of its sub-actions, their object binding, and any constraints. Activities are a partial specification of a scenario: determining their abstract intent and leaving other details open.

Actions prescribe their possible interactions with other actions indirectly, by using flow objects (see 9.3) and resource objects (see 9.4) . *Flow object references* specify the action's inputs and outputs and *resource object references* specify the action's resource claims.

By declaring a reference to an object, an action determines its relation to other actions that reference the very same object without presupposing anything specific about them. For example, one action may reference a data flow object of some type as its input, which another action references as its output. By referencing the same object, the two actions necessarily agree on its properties without having to know about each other. Each action may constrain the attributes of the object. In any consistent scenario, all constraints shall hold; thus, the requirements of both actions are satisfied, as well as any constraints declared in the object itself.

Actions may be *atomic*, in which case their implementation is supplied via one or more **exec body** blocks (see 20.1.2), or they may be *compound*, in which case they contain one or more **activity** statements (see Clause 11) that instantiate and schedule other actions.

An action is declared using the **action** keyword and an *action_identifier*, as shown in Syntax 28.

Copyright © 2026 Accellera. All rights reserved.

9.2.1 Syntax

```
action_declaration ::= action action_identifier [ template_param_decl_list ] [ action_super_spec ]
  { { action_body_item } }
abstract_action_declaration ::= abstract action_declaration
override_action_declaration ::= override action action_identifier { { action_body_item } }
action_super_spec ::= : type_identifier
action_body_item ::=
    activity_declaration
  | override_declaration
  | constraint_declaration
  | action_field_declaration
  | symbol_declaration
  | covergroup_declaration
  | generic_constraint_declaration
  | exec_block_stmt
  | activity_scheduling_constraint
  | attr_group
  | compile_assert_stmt
  | covergroup_instantiation
  | action_body_compile_if
  | annotation_usage action_body_item
  | stmt_terminator
```

Syntax 28—Action declaration

An **action** declaration optionally specifies an *action_super_spec*, a previously defined action type from which the new type inherits its members.

The following also apply:

a) The *activity_declaration* and **body** *exec_block_stmt* (see 20.1.2) action body items are mutually exclusive. An atomic action may specify **body** *exec_block_stmt* items; it shall not specify *activity_declaration* items. A compound action, which contains instances of other actions and *activity_declaration* items, shall not specify **body** *exec_block_stmt* items.

b) An *abstract action* may be declared as a template that defines a base set of field attributes and behavior from which other actions may inherit. Non-abstract derived actions may be instantiated like any other action. Abstract actions shall not be instantiated directly.

c) An abstract action may be derived from another abstract action, but not from a non-abstract action.

d) Abstract actions may be extended, but the action remains abstract and may not be instantiated directly.

9.2.2 Actions and inheritance

An **override** action is virtual with respect to the component in which it is declared. The actual implementation is dictated by the component context in which the action is traversed. An overriding action implicitly inherits from the action that it overrides.

Copyright © 2026 Accellera. All rights reserved.

The following also apply:

a) An action can be declared as **override** in a component only if a same-named action has been declared in a base component.

b) If an **override action** has been declared in a component, any component subtype declaring a same-named action shall also declare it as **override**.

c) Template actions shall not be overridden.

```
component base_c {
  action base_a { }
}
component inh1_c : base_c {
  override action base_a { }
}
component inh2_c : inh1_c {
  override action base_a { }
}
```

Example 57—Override specification

Example 57 illustrates the clauses above. Action `inh1_c::base_a` overrides action `base_c::base_a`, illustrating clause `a`. Action `inh2_c::base_a` overrides `inh1_c::base_a`, and is marked *override* to comply with clause `b`.

9.2.3 Examples

Examples of the atomic, compound, and abstract actions are included in this section.

9.2.3.1 Atomic actions

Examples of an *atomic action* declaration are shown in Example 58.

```
action write {
  output data_buf data;
  rand int size;
  //implementation details
  ...
};
```

Example 58—Atomic action

9.2.3.2 Compound actions

Compound actions instantiate other actions within them and use **activity** statements (see Clause 11) to define the relative scheduling of these sub-actions.

Examples of compound action usage are shown in Example 59.

Copyright © 2026 Accellera. All rights reserved.

```
action sub_a {...};

action compound_a {
  sub_a a1, a2;
  activity {
    a1;
    a2;
  }
}
```

Example 59—Compound action

9.2.3.3 Abstract actions

Abstract action types are used to capture common features of different actions, including actions of different components. Abstract actions may not be traversed directly. Rather, they are used through inheritance, as base types for non-abstract action types. Abstract action types may be declared outside the scope of a component, unlike non-abstract actions, which may only be declared in a component scope.

An example of abstract action usage is shown in Example 60. In this example, abstract action `base` is declared outside a component scope, in package `mypkg`, and subsequently extended in the same package. Action `derived` is declared as a non-abstract subtype of action `base`.

```
package mypkg {
  abstract action base {
    rand int i;
    constraint i>5 && i<10;
  }

// action base remains abstract
  extend action base {
    rand int j;
  }
}

component pss_top {
  import mypkg::*;

  action derived : base {
    constraint i>6;
    constraint j>9;
  }
}
```

Example 60—Abstract action

Copyright © 2026 Accellera. All rights reserved.

9.2.3.4 Overridden actions

```
component TrafficGen {
  action SendTraffic
    rand int in [1..64] n_bursts;
  }
}
component PCIe : TrafficGen {
  override action SendTraffic { }
}
component USB : TrafficGen {
  override action SendTraffic { }
}
component pss_top {
  PCIe  pc1;
  USB   us1;
  action Entry {
    activity {
      do TrafficGen::SendTraffic;
      do PCIe::SendTraffic with { n_bursts > 4; };
    }
  }
}
```

Example 61—Overridden action

The example above shows a base `TrafficGen` component with a `SendTraffic` action and two derivatives: PCIe and USB. The `pss_top` component declares an action that traverses two actions: `TrafficGen::SendTraffic` and `PCIe::SendTraffic`.

Because `TrafficGen::SendTraffic` is declared in the `TrafficGen` component, it can be associated with either the PCIe or the USB component. If it is associated with the PCIe component, the `PCIe::SendTraffic` action will be traversed. If it is associated with the USB component, the `USB::SendTraffic` action will be traversed. This allows `TrafficGen::SendTraffic` to be used as a common interface to component-specific functionality implemented by the PCIe and USB components.

`PCIe::SendTraffic` overrides `TrafficGen::SendTraffic` and, consequently, implicitly inherits from it. Note that the traversal of `PCIe::SendTraffic` constrains the *n_bursts* attribute from the base action. In addition, it can only be associated with the PCIe component because it is declared in the PCIe component.

9.3 Flow objects

A *flow object* represents incoming or outgoing data/control flow for actions, or their pre-condition and post-condition. A flow object can have two modes of reference by actions: **input** and **output**.

9.3.1 Buffer objects

Buffer objects represent data items in some persistent storage that can be written and read. Once their writing is completed, they can be read as needed. Typically, buffer objects represent data or control buffers in internal or external memories. See Syntax 29.

Copyright © 2026 Accellera. All rights reserved.

9.3.1.1 Syntax

```
buffer identifier [ template_param_decl_list ] [ struct_super_spec ] { { struct_body_item } }
```

Syntax 29—Buffer declaration

The following also apply:

a) Note that the buffer type does not imply any specific layout in memory for the specific data being stored.

b) Buffer types can inherit from previously defined structs or buffers.

c) Buffer object reference fields can be declared under actions using the **input** or **output** modifier (see 9.3.4). Instance fields of buffer type (taken as a plain-data type) can only be declared under higher-level buffer types, as their data attribute.

d) A buffer object shall be the output of exactly one action. A buffer object may be the input of any number (zero or more) of actions.

e) Execution of a consuming action that inputs a buffer shall not begin until after the execution of the producing action completes (see Figure 2).

f) An action may not have the same buffer object declared as both an input and an output.

9.3.1.2 Examples

Examples of buffer objects are show in Example 62.

```
struct mem_segment_s {...};
buffer data_buff_s {
    rand mem_segment_s seg;
};
```

Example 62—Buffer object

9.3.2 Stream objects

Stream objects represent transient data or control exchanged between actions during concurrent activity, e.g., over a bus or network, or across interfaces. They represent data item flow or message/notification exchange. See Syntax 30.

9.3.2.1 Syntax

```
stream identifier [ template_param_decl_list ] [ struct_super_spec ] { { struct_body_item } }
```

Syntax 30—Stream declaration

The following also apply:

a) Stream types can inherit from previously defined structs or streams.

b) Stream object reference fields can be declared under actions using the **input** or **output** modifier (see 9.3.4). Instance fields of stream type (taken as a plain-data type) can only be declared under higher-level stream types, as their data attribute.

c) A stream object shall be the output of exactly one action and the input of exactly one action.

Copyright © 2026 Accellera. All rights reserved.

d) The outputting and inputting actions shall begin their execution at the same time, after the same preceding action(s) completes. The outputting and inputting actions are said to run *in parallel*. The semantics of parallel execution are discussed further in 11.3.4.

9.3.2.2 Examples

Examples of stream objects are show in Example 63.

```
struct mem_segment_s {...};
stream data_stream_s {
   rand mem_segment_s seg;
};
```

Example 63—Stream object

9.3.3 State objects

State objects represent the state of some entity in the execution environment at a given time. See Syntax 31.

9.3.3.1 Syntax

```
state identifier [ template_param_decl_list ] [ struct_super_spec ] { { struct_body_item } }
```

Syntax 31—State declaration

The following also apply:

a) The writing and reading of states in a scenario is deterministic. With respect to a pool of state objects, writing shall not take place concurrently to either writing or reading.

b) The initial state of a given type is represented by the built-in Boolean **`initial`** attribute. See 12.5 for more on state pools (and **`initial`**).

c) State object reference fields can be declared under actions using the **input** or **output** modifier (see 9.3.4). Instance fields of state type (taken as a plain-data type) can only be declared under higher-level state types, as their data attribute. It shall be illegal to access the built-in attribute **`initial`** on an instance field.

d) State types can inherit from previously defined structs or states.

e) An action that has an input or output of state object type operates on a pool of the corresponding state object type to which its field is bound. Static pool **bind** directives are used to associate the action with the appropriate state object pool (see 12.3).

f) At any given time, a pool of state object type contains a single state object. This object reflects the last state specified by the output of an action bound to the pool. Prior to execution of the first action that outputs to the pool, the object reflects the initial state specified by constraints involving the **`initial`** built-in field of state object types.

g) The built-in variable `prev` is a reference from this state object to the previous one in the pool. `prev` is a reference to the same type as this state object. The value of `prev` shall be unresolved in the context of the initial state object. `prev` shall only be available within a state type declaration or extension, in relation to this state object itself.

h) An action that inputs a state object reads the current state object from the state object pool to which it is bound.

i) An action that outputs a state object writes to the state object pool to which it is bound, updating the state object in the pool.

Copyright © 2026 Accellera. All rights reserved.

j) Execution of an action that outputs a state object shall complete at any time before the execution of any inputting action begins.

k) Execution of an action that outputs a state object to a pool shall not be concurrent with the execution of any other action that either outputs or inputs a state object from that pool.

l) Execution of an action that inputs a state object from a pool may be concurrent with the execution of any other action(s) that input a state object from the same pool, but shall not be concurrent with the execution of any other action that outputs a state object to the same pool.

9.3.3.2 Examples

Examples of state objects are shown in Example 64.

```
enum mode_e {...};
state config_s {
    rand mode_e mode;
    ...
};
```

Example 64—State object

9.3.4 Using flow objects

Flow object references are specified by actions as **inputs** or **outputs**. These references are used to specify rules for combining actions in legal scenarios. An action that outputs a flow object is said to *produce* that object and an action that inputs a flow object is said to *consume* the object. See Syntax 32.

A consumer may consume flow objects that are produced by multiple producers, and vice versa.

An action can produce or consume a fixed-size array, possibly multidimensional, of flow objects. Declaring such an array is equivalent to declaring multiple distinct object reference fields of the same type.

Copyright © 2026 Accellera. All rights reserved.

9.3.4.1 Syntax

```
action_field_declaration ::=
    attr_field
  | activity_data_field
  | action_handle_declaration
  | object_ref_field_declaration
object_ref_field_declaration ::=
    flow_ref_field_declaration
  | resource_ref_field_declaration
flow_ref_field_declaration ::=
    ( input | output ) flow_object_type object_ref_field { , object_ref_field } ;
flow_object_type ::=
    buffer_type_identifier
  | state_type_identifier
  | stream_type_identifier
object_ref_field ::= identifier { array_dim }
array_dim ::= [ constant_expression ]
```

Syntax 32—Flow object reference

The following apply for arrays of flow object references:

a) Individual elements in the array may be referenced by using the array name and the element index, or indexes in the case of a multidimensional array, in square brackets.

b) A flow object array is specified as entirely input or entirely output. The mode cannot be specified separately for an individual element of the array.

c) The different elements in an array may be bound to different pools. Explicit binding shall be used for array elements associated with different pools. Default (type-based) pool binding applies to all elements of an object-reference array, and therefore cannot be used for this purpose (see 12.3 for more details).

d) For an array of state object references, each object reference shall be bound to a different state pool, since a state pool can store only one state object at a time (see 9.3.3.1 and Example 133).

Copyright © 2026 Accellera. All rights reserved.

9.3.4.2 Examples

Examples of using buffer flow objects are shown in Example 58.

```
struct mem_segment_s {...};
buffer data_buff_s {
    rand mem_segment_s seg;
};
action cons_mem_a {
    input  data_buff_s in_data;
};
action prod_mem_a {
    output data_buff_s out_data;
};
```

Example 65—Buffer flow object

For a timing diagram showing the relative execution of two actions sharing a buffer object, see Figure 2.

Examples of using stream flow objects are shown in Example 66.

```
struct mem_segment_s {...};
stream data_stream_s {
  rand mem_segment_s seg;
};
action cons_mem_a {
  input  data_stream_s in_data;
};
action prod_mem_a {
  output data_stream_s out_data;
};
```

Example 66—Stream flow object

For a timing diagram showing the relative execution of two actions sharing a stream object, see Figure 3.

In Example 67, four **buffer** objects are produced, one by **action** `prod_1b` and three by **action** `prod_3b`, and five **buffer** objects are consumed, one by `cons_1b`, two by `cons_2b_0`, and two by `cons_2b_1`. All the **buffer** objects are produced and consumed from the same **pool,** `buff_p`. All the **buffer** objects have a random integer attribute, `int_attr`. Consumer objects in `cons_2b_0` constrain their `int_attr` attribute to `3`, while in `cons_2b_1`, the first consumer object's `int_attr` attribute is constrained to be greater than or equal to `2`, and the second is constrained to be less than `3`. `prod_3b`'s producer objects `int_attr` attributes are all constrained to `3`.

There is an explicit **bind** to bind the second consumer object in `cons_2b_1` with the first producer object in `prod_3b`, The explicit **bind** constraint will fail because `int_attr` in the consumer object is constrained to be less than `3`, while `int_attr` in the producer object is constrained to `3`. If we remove the explicit **bind**, then that same consumer object will bind to the producer `prod_1b`'s output object because its `int_attr` is constrained to be less than `3`.

Copyright © 2026 Accellera. All rights reserved.

.

```
buffer data_buff {
    rand int int_attr;
};

component flow_object_array_c {
    pool data_buff buff_p;
    bind buff_p *;

    action prod_buff_a {
        output data_buff out_1_buff;
    };

    action prod_3_buff_a {
        output data_buff out_3_buff [3];
    };

    action cons_buff_a {
        input data_buff in_1_buff;
    };

    action cons_2_buff_a {
        input data_buff in_2_buff [2];
    };

    action activity_a {
        prod_buff_a   prod_1b;
        prod_3_buff_a prod_3b;

        cons_buff_a   cons_1b;
        cons_2_buff_a cons_2b_0;
        cons_2_buff_a cons_2b_1;

        activity {
            prod_1b with {out_1_buff.int_attr == 1;};
            prod_3b with {
                foreach (b:out_3_buff) { b.int_attr == 3;};
            };

            cons_1b with { in_1_buff.int_attr == 3;};

            cons_2b_0;
            constraint { foreach (b: cons_2b_0.in_2_buff) {
                b.int_attr == 3;
            };};

            cons_2b_1 with {
                in_2_buff[0].int_attr >= 2 && in_2_buff[1].int_attr < 3;};
            bind cons_2b_1.in_2_buff[1] prod_3b.out_3_buff[0]; // conflict
        };
    };
};
```

Example 67—Multiple producers/consumers using the same buffer pool

An example of use of an array of state object references can be seen in Example 133.

Copyright © 2026 Accellera. All rights reserved.

9.4 Resource objects

Resource objects represent computational resources available in the execution environment that may be assigned to actions for the duration of their execution.

9.4.1 Declaring resource objects

Resource types can inherit from previously defined structs or resources. See Syntax 33. Resources reside in *pools* (see Clause 12) and may be claimed by specific actions.

9.4.1.1 Syntax

```
resource identifier [ template_param_decl_list ] [ struct_super_spec ] { { struct_body_item } }
```

Syntax 33—Resource declaration

The following also apply:

a) Resources have a built-in non-negative integer attribute called **`instance_id`**. This attribute represents the relative index of the resource instance in the pool. The value of **`instance_id`** ranges from **`0`** to *pool_size* **`- 1`**. See also 12.4.

b) There can only be one resource object per **`instance_id`** value for a given pool. Thus, actions referencing a resource object of some type with the same **`instance_id`** are necessarily referencing the very same object and agreeing on all its properties.

c) *Resource object reference* fields can be declared under actions using the **lock** or **share** modifier (see 9.4.2). Instance fields of resource type (taken as a plain-data type) can only be declared under higher-level resource types, as their data attribute.

9.4.1.2 Examples

For examples of how to declare a resource, see Example 58.

```
resource DMA_channel_s {
  rand bit[4] priority;
};
```

Example 68—Declaring a resource

9.4.2 Claiming resource objects

Resource objects may be *locked* or *shared* by actions. This is expressed by declaring the resource reference field of an action. See Syntax 34.

An action can claim a fixed-size array, possibly multidimensional, of resource objects. Declaring such anarray is equivalent to declaring multiple distinct object reference fields of the same type.

Copyright © 2026 Accellera. All rights reserved.

9.4.2.1 Syntax

```
action_field_declaration ::=
    attr_field
  | activity_data_field
  | action_handle_declaration
  | object_ref_field_declaration
object_ref_field_declaration ::=
    flow_ref_field_declaration
  | resource_ref_field_declaration
resource_ref_field_declaration ::=
    ( lock | share) resource_object_type object_ref_field { , object_ref_field } ;
resource_object_type ::= resource_type_identifier
object_ref_field ::= identifier { array_dim }
array_dim ::= [ constant_expression ]
```

Syntax 34—Resource object reference

lock and **share** are modes of resource use by an action. They serve to declare resource requirements of the action and restrict legal scheduling relative to other actions. *Locking* excludes the use of the resource instance by another action throughout the execution of the locking action and *sharing* guarantees that the resource is not locked by another action during its execution.

In a PSS-generated test scenario, no two actions may be assigned the same resource instance if they overlap in execution time and at least one is locking the resource. In other words, there is a strict scheduling dependency between an action referencing a resource object in **lock** mode and all other actions referencing the same resource object instance.

The following apply for arrays of resource object references:

a) Individual elements in the array may be referenced by using the array name and the element index, or indexes in the case of a multidimensional array, in square brackets.

b) A resource object array is specified as entirely locked or entirely shared. The mode cannot be specified separately for an individual element of the array.

c) All elements of a resource object array, from all dimensions, shall be bound to the same pool.

d) When claiming an array of resource objects, the pool size shall be at least as large as the array, with all dimensions, in order to accommodate all distinct resource claims.

9.4.2.2 Examples

Example 62 demonstrates resource claims in lock and share mode. Action `two_chan_transfer` claims exclusive access to two different `DMA_channel_s` instances. It also claims one `CPU_core_s` instance in non-exclusive share mode. While `two_chan_transfer` executes, no other action may claim either instance of the `DMA_channel_s` resource, nor may any other action lock the `CPU_core_s` resource instance.

Copyright © 2026 Accellera. All rights reserved.

```
resource DMA_channel_s {
  rand bit[4] priority;
};
resource CPU_core_s {...};
action two_chan_transfer {
  lock DMA_channel_s chan_A;
  lock DMA_channel_s chan_B;
  share CPU_core_s ctrl_core;
...
};
```

Example 69—Resource object

In Example 70, there is a pool of 16 resource objects of type `config`. The **action** `baz_lock_a` claims a lock for 8 resource objects. The **action** `baz_share_a` claims to share 16 resource objects. The **action** `entry_a` can legally traverse two `baz_share_a` actions in parallel, as the same resource object can be shared between concurrent activities. It can also legally traverse two `baz_lock_a` actions in parallel because overall there are 16 resource objects and each action instance consumes only 8.

```
resource config {}

component foo_c {
  pool[16] config config_p;
  bind config_p *;

  action baz_lock_a {
    lock config  config_object[8];
  }

  action baz_share_a {
    share config config_object[16];
  }

  action entry_a {
    activity {
      parallel {
        do baz_share_a;
        do baz_share_a;
      }
      parallel {
        do baz_lock_a;
        do baz_lock_a;
      }
    }
  }
}
```

Example 70—Locking and sharing arrays of resource objects

Copyright © 2026 Accellera. All rights reserved.

10. Template types

10.1 General

Template types in PSS provide a way to define generic parameterized types. A template type is either a **struct** type or a modeling-element type.

In many cases, it is useful to define a generic parameterizable type that can be instantiated with different parameter values (e.g., array sizes or data types). Template types maximize reuse, avoid writing similar code for each parameter value (value or data type) combination, and allow a single specification to be used in multiple places.

Template types shall be explicitly instantiated by the user, and only an explicit instantiation of a template type represents an actual type.

The following sections describe how to define, use, and extend a template type when using the PSS input.

10.2 Template type declarations

A **struct** type or a modeling-element type whose declaration includes a template parameter list is a template type.

A *template type* declaration specifies a list of formal type or value template parameter declarations. The parameters are provided as a comma-separated list enclosed in angle brackets (<>) following the name of the template type.

A template type may inherit from another template or non-template data type. A non-template type may inherit from a template type instance. In both cases, the same inheritance rules and restrictions as for the corresponding non-template type of the same type category are applied (e.g., a template **struct** may inherit from a **struct**, or from a template **struct**).

The declaration syntax for **struct**s and modeling elements includes an optional *template_param_decl_list* nonterminal. Presence of this nonterminal in a **struct** or modeling element declaration denotes that the declared type is a template generic type.

10.2.1 Syntax

```
struct_declaration ::= struct_kind identifier [ template_param_decl_list ] [ struct_super_spec ] { { struct_body_item } }
component_declaration ::= component component_identifier [ template_param_decl_list ] [ component_super_spec ] { { component_body_item } }
action_declaration ::= action action_identifier [ template_param_decl_list ] [ action_super_spec ] { { action_body_item } }
monitor_declaration ::= monitor monitor_identifier [ template_param_decl_list ] [ monitor_super_spec ] { { monitor_body_item } }
template_param_decl_list ::= < template_param_decl { , template_param_decl } >
template_param_decl ::= type_param_decl | value_param_decl
```

Syntax 35—Template type declaration

Copyright © 2026 Accellera. All rights reserved.

10.2.2 Example

Generic template-type declaration for various type categories are shown in Example 71.

```
struct my_template_s <type T> {
  T  t_attr;
}

buffer my_buff_s <type T> {
  T  t_attr;
}

abstract action my_consumer_action <int width, bool is_wide> {
  compile assert (width > 0);
}
monitor my_sequence <monitor m1, monitor m2> {
  activity {
    do m1; do m2;
  }
}


component eth_controller_c <struct ifg_config_s, bool full_duplex = true> {
}
```

Example 71—Template type declarations

10.3 Template parameter declarations

A template parameter is declared as either a type or a value parameter. All template parameters have a name and an optional default value. All parameters subsequent to the first one that is given a default value shall also be given default values. Therefore, the parameters with defaults shall appear at the end of the parameter list. Specifying a parameter with a default value followed by a parameter without a default value shall be reported as an error.

A template parameter can be referenced using its name inside the body and the supertype specification of the template type and all subsequent generic template type extensions, including the template type instance extensions. A template parameter may not be referenced from within subtypes that inherit from the template type that originally defined the parameter.

10.3.1 Template value parameter declarations

Value parameters are given a data type and optionally a default value, as shown below.

10.3.1.1 Syntax

```
value_param_decl ::= data_type identifier [ = constant_expression ]
```

Syntax 36—Template value parameter declaration

The following also apply:

a) A value parameter can be referenced using its name anywhere a constant expression is allowed or expected inside the body and the supertype specification of the template type.

Copyright © 2026 Accellera. All rights reserved.

b) Valid data types for a *value_param_decl* are the scalar types, except **chandle**.

c) The default value, if provided, may also reference one or more of the previously defined parameters.

d) To avoid parsing ambiguity, a Boolean greater-than (>) or less-than (<) expression provided as a default value shall be enclosed in parentheses.

10.3.1.2 Examples

An example of declaring an action type that consumes a varying number of resources is shown in Example 72.

```
action my_consumer_action <int n_locks = 4> {
  compile assert (n_locks in [1..16]);
  lock my_resource res[n_locks];
}
```

Example 72—Template value parameter declaration

Example 73 contains a Boolean greater-than expression that shall be enclosed in parentheses and depends on a previous parameter:

```
action my_consumer_action <int width, bool is_wide = (width > 10) > {
  compile assert (width > 0);
}
```

Example 73—Another template value parameter declaration

10.3.2 Template type parameter declarations

Type parameters are prefixed with either the **type** keyword or a type-category keyword in order to identify them as type parameters.

When the **type** keyword is used, the parameter is fully generic. In other words, it can take on any type.

Specifying category type parameters provides more information to users of a template type on acceptable usage and allows tools to flag usage errors earlier. A category type parameter enforces that a template instance parameter value shall be of a certain category/class of type (e.g., **struct**, **action**, **numeric**, etc.). A category type parameter (other than *numeric*) can be further restricted such that the specializing type (the parameter value provided on instantiation) shall be related via inheritance to a specified base type.The syntax for declaring a type parameter is shown below.

```
type_param_decl ::= generic_type_param_decl | category_type_param_decl
generic_type_param_decl ::= type identifier [ = type_identifier ]
category_type_param_decl ::= type_category identifier [ type_restriction ] [ = type_identifier ]
type_restriction ::= : type_identifier
type_category ::=
    ref_type_category
  | plain_type_category
```

Syntax 37—Template type parameter declaration

Copyright © 2026 Accellera. All rights reserved.

The following also apply:

a) A type parameter can be referenced using its name anywhere inside the body and the supertype specification of the template type where a type is allowed or expected.

b) The default value, if provided, may also reference one or more of the previously defined parameters.

10.3.2.1 Examples

Examples of a generic type and a category type parameter are shown in Example 74.

```
struct my_container_s <struct T> {
  T t_attr;
}

struct my_template_s <type T> {
  T t_attr;
}
```

Example 74—Template generic type and category type parameters

In the example above, the template parameter `T` of `my_container_s` shall be of **struct** type, while in the case of `my_template_s`, the template parameter `T` may take on any type.

An example of how to use type restrictions in the case of a type-category parameter is shown in Example 75.

```
struct base_t {
  rand bit[4] core;
}

struct my_sub1_t : base_t {
  rand bit[4] add1;
}

struct my_sub2_t : base_t {
  rand bit[4] add2;
}

buffer b1 : base_t { }
buffer b2 : base_t { }

abstract action my_action_a <buffer B : base_t> {
}

struct my_container_s  <struct T : base_t = my_sub1_t> {
  T t_attr;
  constraint t_attr.core >= 1;
}
```

Example 75—Template parameter type restriction

In the example above, the template parameter **`T`** of `my_container_s` shall be of type `base_t` or one of its **struct** subtypes (`my_sub1_t` or `my_sub2_t`, but not `b1` or `b2`). This allows `my_container_s` to reasonably assume that **`T`** contains an attribute named '`core`', and communicates this requirement to users of this type and to the PSS processing tool. The template parameter `B` of `my_action_a` shall be of one of the **buffer** subtypes of `base_t` (`b1` or `b2`).

Copyright © 2026 Accellera. All rights reserved.

The base type of the template type may also be a type parameter. In this way, the inheritance can be controlled when the template type is instantiated.

In Example 76, the `my_container_s` template **struct** inherits from the **struct** type template type parameter.

```
struct my_base1_t {
  rand int attr1;
}

struct my_base2_t {
  rand int attr2;
}

struct my_container_s <struct T> : T {
}

struct top_s {
  rand my_container_s <my_base1_t> cont1;
  rand my_container_s <my_base2_t> cont2;
  constraint cont1.attr1  == cont2.attr2;
}
```

Example 76—Template parameter used as base type

10.4 Template type instantiation

A template type is instantiated using the name of the template type followed by the parameter value list (specialization) enclosed in angle brackets (<>). Template parameter values are specified positionally.

The explicit instantiation of a template type represents an actual type. All explicit instantiations provided with the same set of parameter values are the same actual type.

10.4.1 Syntax

```
type_identifier ::= [ :: ] type_identifier_elem { :: type_identifier_elem }
type_identifier_elem ::= identifier [ template_param_value_list ]
template_param_value_list ::= < [ template_param_value { , template_param_value } ] >
template_param_value ::= constant_expression | data_type
```

Syntax 38—Template type instantiation

The following also apply:

a) Parameter values shall be specified for all parameters that were not given a default value.

b) An instance of a template type shall always specify the angle brackets (<>), even if no parameter value overrides are provided for the defaults.

c) The specified parameter values shall comply with parameter categories and parameter type restrictions specified for each parameter in the original template declaration, or an error shall be generated.

d) To avoid parsing ambiguity, a Boolean greater-than (>) or less-than (<) expression provided as a parameter value shall be enclosed in parentheses.

Copyright © 2026 Accellera. All rights reserved.

10.4.2 Examples

```
struct base_t {
  rand bit[4] core;
}

struct my_sub1_t : base_t {
  rand bit[4] add1;
}

struct my_sub2_t : base_t {
  rand bit[4] add2;
}

struct my_container_s <struct T : base_t = my_sub1_t> {
  T t_attr;
  constraint t_attr.core >= 1;
}

struct top_s {
  my_container_s<>          my_sub1_container_attr;
  my_container_s<my_sub2_t> my_sub2_container_attr;
}
```

Example 77—Template type instantiation

In Example 77 above, two attributes of `my_container_s` type are created. The first uses the default parameter value. The second specifies the `my_sub2_t` type as the value for the **`T`** parameter.

Type qualification for an action declared in a template component is shown in Example 78 below.

```
component my_comp1_c <int bus_width = 32> {
  action my_action1_a { }
  action my_action2_a <int nof_iter = 4> { }
}

component pss_top {
  my_comp1_c<64> comp1;
  my_comp1_c<32> comp2;

  action test {
    activity {
      do my_comp1_c<64>::my_action1_a;
      do my_comp1_c<64>::my_action2_a<>;
      do my_comp1_c::my_action1_a;  // Error -  my_comp1_c must be specialized
      do my_comp1_c<>::my_action1_a;
    }
  }
}
```

Example 78—Template type qualification

Copyright © 2026 Accellera. All rights reserved.

Example 79 depicts various ways of overriding the default values. In the example below, the `my_struct_t<2>` instance overrides the parameter `A` with `2`, and preserves the default values for parameters `B` and `C`. The `my_struct_t<2, 8>` instance overrides the parameter `A` with `2`, parameter `B` with `8`, and preserves the default value for `C`.

```
struct my_s_1 { }
struct my_s_2 { }

struct my_struct_t <int A = 4, int B = 7, int C = 3> {  }

struct container_t {
  my_struct_t<2>   a; // instantiated with <2, 7, 3>
  my_struct_t<2,8> b; // instantiated with <2, 8, 3>
}
```

Example 79—Overriding the default values

10.5 Template type user restrictions

A generic template type may not be used in the following contexts:

- As a root component
- As a root action
- As an inferred action to complete a partially specified scenario

Template types are explicitly instantiated by the user, and only an explicit instantiation of a template type represents an actual type. Only action actual types can be inferred to complete a partially specified scenario. The root component and the root action shall be actual types.

Template types may not be used as parameter types or return types of imported and exported functions.

Copyright © 2026 Accellera. All rights reserved.

11. Action activities

When a *compound action* includes multiple operations, these behaviors are described within the **action** using one or more **activity** statements. An *activity* specifies the set of actions to be executed and the scheduling relationship(s) between them. If more than one activity is specified in an action, the execution semantics are the same as if the activity statements were combined in a **schedule** statement (see 11.3.5 and 11.6). A reference to an action within an activity is via an *action handle*, and the resulting *action traversal* causes the referenced action to be evaluated and randomized (see 11.3.1).

An activity, on its own, does not introduce any scheduling dependencies for its containing action. However, flow object or resource scheduling constraints of the sub-actions may introduce scheduling dependencies for the containing action relative to other actions in the system.

11.1 Activity declarations

Because activities are explicitly specified as part of an action, activities themselves do not have a separate name. Relative to the sub-actions referred to in the activity, the action that contains the activity is referred to as the *context action*.

11.2 Activity constructs

Each node of an activity represents an action, with the activity specifying the temporal, control, and/or data flow between them. These relationships are described via activity rules, which are explained herein. See also Syntax 39.

Copyright © 2026 Accellera. All rights reserved.

11.2.1 Syntax

```
activity_declaration ::= activity { { activity_stmt } }
activity_stmt ::=
    [ label_identifier :] labeled_activity_stmt
  | activity_action_traversal_stmt
  | activity_data_field
  | activity_bind_stmt
  | action_handle_declaration
  | activity_constraint_stmt
  | activity_scheduling_constraint
  | activity_super_stmt
  | annotation_usage activity_stmt
  | stmt_terminator
labeled_activity_stmt ::=
    activity_sequence_block_stmt
  | activity_parallel_stmt
  | activity_schedule_stmt
  | activity_repeat_stmt
  | activity_foreach_stmt
  | activity_select_stmt
  | activity_if_else_stmt
  | activity_match_stmt
  | activity_replicate_stmt
  | activity_atomic_block_stmt
  | symbol_call
```

Syntax 39—Activity statement

11.3 Action scheduling statements

By default, statements in an activity specify sequential behaviors, subject to data flow constraints. In addition, there are several statements that allow additional scheduling semantics to be specified. Statements within an activity may be nested, so each element within an activity statement is referred to as a sub-activity.

11.3.1 Action traversal statement

An *action traversal statement* designates the point in the execution of an activity where an action is randomized and evaluated (see Syntax 40). The action being traversed may be specified via an action handle referring to an action field or local variable that was previously declared. Alternatively, the action being traversed may be specified by type, in which case a label, if specified, serves as an action handle. In the absence of a label, the action instance is anonymous.

Copyright © 2026 Accellera. All rights reserved.

11.3.1.1 Syntax

```
activity_action_traversal_stmt ::=
  action_handle_traversal_stmt
  | action_type_traversal_stmt

action_handle_traversal_stmt ::=
    identifier { [ expression ] } [ action_initializer_list ]
    inline_constraints_or_empty

action_type_traversal_stmt ::=
  [ label_identifier : ] do type_identifier [ action_initializer_list ]
    inline_constraints_or_empty

inline_constraints_or_empty ::=
    action_activity_inline_constraints_or_empty
  | monitor_activity_inline_constraints_or_empty
action_activity_inline_constraints_or_empty ::=
    with constraint_set
  | ;
```

Syntax 40—Action traversal statement

identifier names a unique action handle or variable in the context of the containing action type or activity scope. If *identifier* refers to an *action handle array* (see 11.3.2), then a specific array element or (if the array is multidimensional) a sub-array may be specified with one or more optional array subscripts. The alternative forms are specified by the keyword **do**, followed by an action-type specifier. Given a *label_identifier*, the action instance can be referenced using the label. In the absence of a *label_identifier*, the action instance is anonymous. Either form of the action traversal statement may include an optional in-line constraint that shall be specified using *action_activity_inline_constraints_or_empty*.

A list of initialization assignments may be associated with the action traversal, as well as with an action-handle declaration. If specified, these initializations are applied when the action is traversed. If specified, initializations associated with the action-handle declaration are applied first, followed by initializations associated with the action traversal statement. Initialization-assignment patterns are a list of assignments to perform. Unlike struct literals (which have a similar syntactic form), initialization assignment patterns can refer to hierarchical paths within the action handle to which they are applied.

The following also apply:

a) An action handle is considered uninitialized until it is first traversed. The fields within the **action** cannot be referenced in an **exec** block or conditional activity statement until after the action is first traversed.

b) Upon entry to an **activity** scope, all action handles traversed in that scope are reset to an uninitialized state.

c) The labeled traversal statement is semantically equivalent to a traversal statement with an explicitly declared action variable. With this form, the *label_identifier* serves as an action handle, equivalent to an explicitly declared variable of the specified action type in the enclosing activity scope.

Copyright © 2026 Accellera. All rights reserved.

d) The *anonymous action traversal* statement is semantically equivalent to the other two forms with the exception that it does not create an action handle that may be referenced from elsewhere in the stimulus model.

The following also apply for action traversal statements in action activity only:

a) The action variable is randomized and evaluated at the point in the flow where the statement occurs. The variable may be of an action type or a data type declared in the context action with the **action** modifier. In the latter case, it is randomized, but has no observed execution or duration (see Example 173).

b) The steps that occur as part of the action traversal are as follows:
   i) Initialization assignments from the initialization list associated with the action handle (if present) are applied.
   ii) Initialization assignments from the initialization list associated with the action traversal (if present) are applied.
   iii) The **pre_solve** *exec block* (if present) is executed.
   iv) Random values are selected for **rand** fields.
   v) The **post_solve** *exec block* (if present) is executed.
   vi) Resolve memory addresses for address claims in action.
   vii) The **pre_body** *exec block* (if present) is executed.
   viii) The **body** *exec block* (if present) is executed.
   ix) The **activity** block (if present) is evaluated.
   x) The validity of the constraint system is confirmed, given any changes by the **post_solve** or **body** *exec blocks*.

c) A named action handle may only be traversed once in the following scopes and nested scopes thereof:
   1) sequential activity scope (e.g., **sequence** or **repeat**)
   2) **parallel**
   3) **schedule**

d) The identifier on the left side of the initialization shall be an attribute of the traversed action type. Other attribute field paths are resolved and evaluated using the same resolution rules as for in-line constraints (13.1.4)

e) Other aspects that impact action-evaluation scheduling, are covered via binding inputs or outputs (see 9.4), resource claims (see 9.4.2), or attribute value assignment.

f) When an action is traversed, its component context will be randomly chosen from the instantiated components of the correct type in the component subtree, starting with the component context of the action containing the activity in which it is traversed.

11.3.1.2 Action attribute initialization

The initial value of action attributes can be specified as part of the action-traversal statement. This enables data to be passed from an activity to a sub-action traversed within the activity.

The following also apply:

a) The evaluation order of assignments within a given initialization list is not specified.

b) Initialization statements may not mutate action-level attributes.

c) To ensure predictable behavior, functions called as part of an initialization expression shall not have side effects.

d) The identifier on the left side of the initialization shall be an attribute of the traversed action type.

Copyright © 2026 Accellera. All rights reserved.

e) Only fields of the traversed action that are accessible in **pre_solve** can be initialized. See 13.4.12,

f) Only fields of the parent action that are accessible in **post_solve** can be referred to in the assignment.

11.3.1.3 Examples

Example 80 shows an example of traversing an action handle. Action `A` is an atomic action that contains a 4-bit random field `f1`. Action `B` is a compound action encapsulating an activity involving two invocations of action `A`. The constraints for `A` apply to the evaluation of `a1` and `a2`. An additional constraint is applied to `a2`, specifying that `f1` shall be less than `10`. Execution of action `B` results in two sequential evaluations of action `A`.

```
action A {
  rand bit[4]   f1;
  constraint f1 > 3;
}

action B {
  A a1, a2;

  activity {
    a1;
    a2 with {
      f1 < 10;
    };
  }
}
```

Example 80—Action traversal

Example 81 shows an example of anonymous action traversal, including in-line constraints.

```
action A {
  rand bit[4]   f1;
  constraint f1 > 3;
}

action B {
  activity {
    do A;
    do A with {f1 < 10;};
  }
}
```

Example 81—Anonymous action traversal

Example 82 shows the use of a label of an action traversal statement to constrain a sub-action instance from a higher activity context.

Copyright © 2026 Accellera. All rights reserved.

```
action mem2mem_chain {
  activity {
    do mem_c::load_buff;
    repeat (10) {
      select {
        xfer: do dma_c::mem2mem_xfer;
        cpy:  do cpu_c::memcpy;
      }
    }
  }
}

action my_test {
  activity {
    do mem2mem_chain with { xfer.size > 10; };
  }
}
```

Example 82—Labeled action traversal

Example 83 shows an example of traversing a compound action as well as a random action variable field. The activity for action `C` traverses the random action variable field `max`, then traverses the action-type field `b1`. Evaluating this activity results in a random value being selected for `max`, then the sub-activity of `b1` being evaluated, with `a1.f1` constrained to be less than or equal to `max`.

```
action A {
  rand bit[4]   f1;
  constraint f1 > 3;
}

action B {
  A a1, a2;

  activity {
    a1;
    a2 with {
      f1 < 10;
    };
  }
}

action C {
  action bit[4] max;
  B b1;

  activity {
    max;
    b1 with {
      a1.f1 <= max;
    };
  }
}
```

Example 83—Compound action traversal

Copyright © 2026 Accellera. All rights reserved.

An initializer list can optionally be associated with an action-handle declaration, an action-traversal statement, or both. When both are supplied, the initialization statements from the list associated with the action-handle declaration are evaluated first, followed by the initialization statements from the list associated with the action traversal.

```
action A {rand bit knob ; bit a, b;}

action test {
    rand bool b;
    A a_h {.a=1, .b=2};

    activity {
        if (b) {
            a_h {.a=3} with {knob == 0};
        } else {
            a_h {.b=1} with {knob == 1};
        }
    }
}
```

Example 84—-Initialization at declaration and traversal levels

In Example 84, an initializer list is associated with the declaration of action handle `a_h` as well as with the two traversals of the action handle. When the handle is traversed, the initialization statements associated with the handle declaration are applied first, followed by the initialization statements associated with the traversal statement.

When the *true* branch of the if/else statement is evaluated (`b==true`):

- `a_h.a=3; a_h.b=2`

When the *false* branch of the if/else statement is evaluated (`b==false`):

- `a_h.a=1 ; a_h.b=1`

11.3.2 Action handle array traversal

Arrays, possibly multidimensional, of action handles may be declared within an action. These *action handle arrays* may be traversed as a whole or traversed as individual elements. In the case of a multidimensional array, a sub-array may also be traversed as a whole.

The semantics of traversing individual action handle array elements are the same as those of traversing individually-declared action handles.

When traversing an array (including a sub-array of a multidimensional array) as a whole, an attribute initialization list or an inline constraint shall not be specified.

Example 85 below shows traversing an individual action handle array element and one action handle. The semantics of both action traversal statements are the same.

Copyright © 2026 Accellera. All rights reserved.

```
component pss_top {
    action A { }
    action entry {
        A   a_arr[4];
        A   a1, a2, a3, a4;
        activity {
            a_arr[0];
            a1;
        }
    }
}
```

Example 85—Individual action handle array element traversal

When an action handle array is traversed as a whole, each array element is traversed independently according to the semantics of the containing scope.

Example 86 below shows an action that traverses the elements of the `a_arr` action handle array in two ways, depending on the value of a **rand** action attribute. Both ways of traversing the elements of `a_arr` have identical semantics.

```
component pss_top {
    action A { }
    action entry {
        rand bool traverse_arr;
        A      a_arr[2];
        activity {
            if (traverse_arr) {
                a_arr;
            } else {
                a_arr[0];
                a_arr[1];
            }
        }
    }
}
```

Example 86—Action handle array traversal

Example 87 below declares a two-dimensional array of action handles. It shows an action that choses between three alternatives: traversing two individual elements, traversing one sub-array, traversing the entire array.

Copyright © 2026 Accellera. All rights reserved.

```
component pss_top {
    action A { }
    action entry {
        rand int in [0..2] traverse_mode;
        A a_arr[3][2];
        activity {
            match (traverse_mode) {
                [0]: {
                    a_arr[2][0];
                    a_arr[1][1];
                }
                [1]: a_arr[1]; // same as:
                               // { a_arr[1][0]; a_arr[1][1]; }
                [2]: a_arr; // same as:
                            // { a_arr[0][0]; a_arr[0][1];
                            //   a_arr[1][0]; a_arr[1][1];
                            //   a_arr[2][0]; a_arr[2][1]; }
            }
        }
    }
}
```

Example 87—Two-dimensional action handle array traversal

The contexts in which action handle arrays may be traversed, and the resulting semantics, are described in the table below.

Table 24—Action handle array traversal contexts and semantics

Context	Semantics
parallel	All array elements are scheduled for traversal in parallel.
schedule	All array elements are scheduled for traversal independently.
select	One array element is randomly selected and traversed. If the array is multidimensional, only one leaf element, i.e., one specific action handle, is selected and traversed.
sequence	All array elements are scheduled for traversal in sequence from 0 to N-1. If the array is multidimensional, each outer dimension is fully traversed with its subdimensions, before traversing the next subdimension starts. For example, given a two-dimensional action handle array declared as `arr[2][2]`, its elements are traversed in this order: `a[0][0]`, `a[0][1]`, `a[1][0]`, `a[1][1]`.

Copyright © 2026 Accellera. All rights reserved.

11.3.3 Sequential block

An *activity sequence block* statement specifies sequential scheduling between sub-activities (see Syntax 41).

11.3.3.1 Syntax

```
activity_sequence_block_stmt ::= [ sequence ] { { activity_stmt } }
```

Syntax 41—Activity sequence block

The following also apply:

a) Statements in a sequential block execute in order so that one sub-activity completes before the next one starts.

b) Formally, a sequential block specifies sequential scheduling between the sets of action executions per the evaluation of $activity_stmt_1$ .. $activity_stmt_n$, keeping all scheduling dependencies within the sets and introducing additional dependencies between them to obtain sequential scheduling (see 6.3.2).

c) Sequential scheduling does not rule out other inferred dependencies affecting the nodes in the sequence block. In particular, there may be cases where additional action executions shall be scheduled in between sub-activities of subsequent statements.

11.3.3.2 Examples

Assume `A` and `B` are **action** types that have no rules or nested activity (see Example 88).

Action `my_test` specifies one execution of action `A` and one of action `B` with the scheduling dependency `(A) -> (B)`; the corresponding observed behavior is `{start A, end A, start B, end B}`.

Now assume action `B` has a state precondition which only action `C` can establish. `C` may execute before, concurrently to, or after `A`, but it shall execute before `B`. In this case the scheduling dependency relation would include `(A) -> (B)` and `(C) -> (B)` and multiple behaviors are possible, such as `{start C, start A, end A, end C, start B, end B}`.

Finally, assume also `C` has a state precondition which only `A` can establish. Dependencies in this case are `(A) -> (B)`, `(A) -> (C)` and `(C) -> (B)` (note that the first pair can be reduced) and, consequently, the only possible behavior is `{start A, end A, start C, end C, start B, end B}`.

```
action my_test {
   A a;
   B b;
   activity {
      a;
      b;
   }
};
```

Example 88—Sequential block

Example 89 shows all variants of specifying sequential behaviors in an activity. By default, statements in an activity execute sequentially. The **sequence** keyword is optional, so placing sub-activities inside braces ({})

Copyright © 2026 Accellera. All rights reserved.

is the same as an explicit **sequence** statement, which includes sub-activities inside braces. The examples show a total of six sequential actions: `A`, `B`, `A`, `B`, `A`, `B`.

```
action my_test {
   A a1, a2, a3;
   B b1, b2, b3;
   activity {
      a1;
      b1;
      {a2; b2;};
      sequence{a3; b3;};
   }
};
```

Example 89—Variants of specifying sequential execution in activity

11.3.4 Parallel

The *parallel statement* specifies sub-activities that execute concurrently (see Syntax 42).

11.3.4.1 Syntax

```
activity_parallel_stmt ::= parallel [ activity_join_spec ] { { activity_stmt } }
```

Syntax 42—Parallel statement

The following also apply:

a) Parallel activities are invoked in a synchronized way and then proceed without further synchronization until their completion. Parallel scheduling guarantees that the invocation of an action in one sub-activity branch does not wait for the completion of any action in another.

b) Formally, the **parallel** statement specifies parallel scheduling between the sets of action executions per the evaluation of *activity_stmt*$_1$ .. *activity_stmt*$_n$, keeping all scheduling dependencies within the sets, ruling out scheduling dependencies across the sets, and introducing additional scheduling dependencies to initial action executions in each of the sets in order to obtain a synchronized start (see 6.3.2).

c) In the absence of an *activity_join_spec* (see 11.3.6), execution of the activity statement following the **parallel** block is scheduled to begin after all parallel branches have completed. When an *activity_join_spec* is specified, execution of the activity statement following the **parallel** block is scheduled based on the *join* specification.

11.3.4.2 Examples

Assume `A`, `B`, and `C` are **action** types that have no rules or nested activity (see Example 90).

The activity in action `my_test` specifies two dependencies `(a) -> (b)` and `(a) -> (c)`. Since the executions of both `b` and `c` have the exact same scheduling dependencies, their invocation is synchronized.

Now assume action type `C` inputs a buffer object and action type `B` outputs the same buffer object type, and the input of `c` is bound to the output of `b`. According to buffer object exchange rules, the inputting action shall be scheduled after the outputting action. But this cannot satisfy the requirement of parallel scheduling, according to which an action in one branch cannot wait for an action in another. Thus, in the presence of a separate scheduling dependency between `b` and `c`, this activity shall be illegal.

Copyright © 2026 Accellera. All rights reserved.

```
action my_test {
   A a;
   B b;
   C c;
   activity {
      a;
      parallel {
         b;
         c;
      }
   }
};
```

Example 90—Parallel statement

In Example 91, the semantics of the **parallel** construct require the sequences `{A,B}` and `{C,D}` to start execution at the same time. The semantics of the sequential block require that the execution of `B` follows `A` and `D` follows `C`. It is illegal to have any scheduling dependencies between sub-activities in a **parallel** statement, so neither `A` nor `B` may have any scheduling dependencies relative to either `C` or `D`.

Even though actions `A` and `D` lock the same resource type from the same pool, the pool contains a sufficient number of resource instances such that there are no scheduling dependencies between the actions. If `pool_R` contained only a single instance, there would be a scheduling dependency in that `A` and `D` could not overlap, which would violate the rules of the **parallel** statement.

```
resource R{...}
pool [4] R R_pool;
bind R_pool *;
action A { lock R r; }
action B {}
action C {}
action D { lock R r; }

action my_test {
   activity {
      parallel {
         {do A; do B;}
         {do C; do D;}
      }
   }
}
```

Example 91—Another parallel statement

11.3.5 Schedule

The **schedule** statement specifies that the PSS processing tool shall select a legal order in which to evaluate the sub-activities, provided that one exists. See Syntax 43.

Copyright © 2026 Accellera. All rights reserved.

11.3.5.1 Syntax

```
activity_schedule_stmt ::= schedule [ activity_join_spec ] { { activity_stmt } }
```

Syntax 43—Schedule statement

The following also apply:

a) All activities inside the **schedule** block shall execute, but the PSS processing tool is free to execute them in any order that satisfies their other scheduling requirements.

b) Formally, the **schedule** statement specifies that any scheduling of the combined sets of action executions per the evaluation of $activity_stmt_1$ .. $activity_stmt_n$ is permissible, as long as it keeps all scheduling dependencies within the sets and introduces (at least) the necessary scheduling dependencies across the sets in order to comply with the rules of input-output binding of actions and resource assignments.

c) In the absence of an *activity_join_spec* (see 11.3.6), execution of the activity statement following the **schedule** block is scheduled to begin after all statements within the block have completed. When an *activity_join_spec* is specified, execution of the activity statement following the **schedule** block is scheduled based on the *join* specification.

11.3.5.2 Examples

Consider the code in Example 92, which is similar to Example 90, but uses a **schedule** block instead of a **parallel** block. In this case, the following executions are valid:

a) The sequence of action nodes `a`, `b`, `c`.

b) The sequence of action nodes `a`, `c`, `b`.

c) The sequence of action node `a`, followed by `b` and `c` run in any order, subject to other scheduling constraints.

```
action my_test {
   A a;
   B b;
   C c;
   activity {
      a;
      schedule {
         b;
         c;
      }
   }
};
```

Example 92—Schedule statement

Note that neither `b` nor `c` may start execution until after the completion of `a`, and the start of execution for either may be subject to additional scheduling constraints. In contrast to `b` and `c` executing in parallel, as in Example 90, there may be scheduling dependencies between `b` and `c` in the **schedule** block. The scheduling graph for the activity is shown here:

Copyright © 2026 Accellera. All rights reserved.

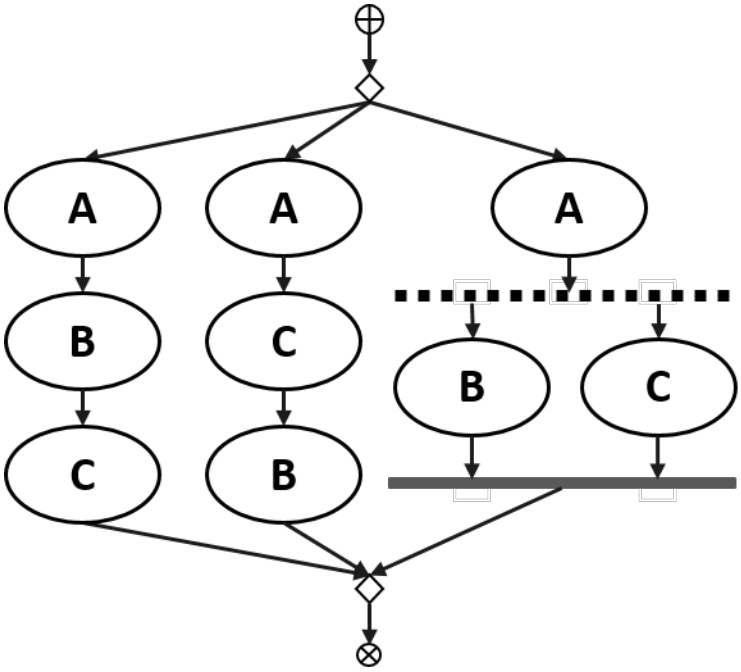


Figure 6—Scheduling graph of activity with schedule block

For the case where `b` and `c` overlap, the runtime behaviors will execute as shown here:

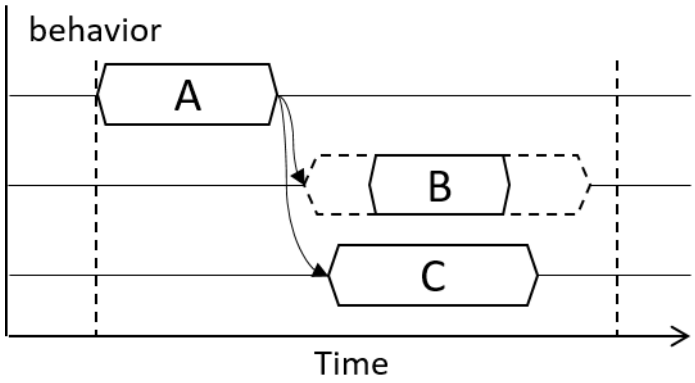


Figure 7—Runtime behavior of activity with schedule block

In contrast, consider the code in Example 93. In this case, any execution order in which both `B` comes after `A` and `D` comes after `C` is valid.

If both `A` and `D` wrote to the same state variable, they would have to execute sequentially. This is in addition to the sequencing of `A` and `B` and of `C` and `D`. In the case where `D` writes before `A`, the sequence would be {`C`, `D`, `A`, `B`}. In the case where `A` writes before `D`, the runtime behavior would be as shown in Figure 8.

Copyright © 2026 Accellera. All rights reserved.

```
action A {}
action B {}
action C {}
action D {}

action my_test {
   activity {
      schedule {
         {do A; do B;}
         {do C; do D;}
      }
   }
}
```

Example 93—Scheduling block with sequential sub-blocks

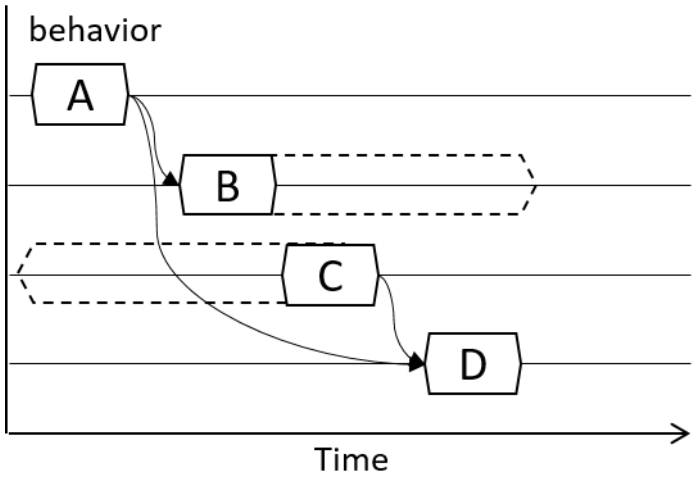


Figure 8—Runtime behavior of scheduling block with sequential sub-blocks

11.3.6 Fine-grained scheduling specifiers

Fine-grained scheduling specifiers modify the termination semantics for **parallel** and **schedule** blocks (see Syntax 42, Syntax 43, and Syntax 44). The semantics of fine-grained scheduling are defined strictly at the activity scheduling level. The semantics do not assume that any runtime execution information is incorporated by the PSS processing tool in the scheduling process. Activity scheduling in the presence of a fine-grained scheduling specifier is still subject to all other scheduling rules.

Copyright © 2026 Accellera. All rights reserved.

11.3.6.1 Syntax

```
activity_join_spec ::=
    activity_join_branch
  | activity_join_select
  | activity_join_none
  | activity_join_first
activity_join_branch ::= join_branch ( label_identifier { , label_identifier } )
activity_join_select ::= join_select ( expression )
activity_join_none ::= join_none
activity_join_first ::= join_first ( expression )
```

Syntax 44—Activity join specification

The following also apply:

a) **join_branch** accepts a list of labels referring to labeled activity statements. The activity statement following the fine-grained scheduling block is scheduled after all the listed activity statements have completed.
   1) The *label_identifier* used in the **join_branch** specification shall be the label of a top-level branch within the **parallel** or **schedule** block to which the **join_branch** specification is applied.
   2) When the *label_identifier* used in the **join_branch** specification applies to traversal of an array, the activity statement following the fine-grained scheduling block is scheduled after all actions in the array have completed.

b) **join_select** accepts an *expression* specifying the number of top-level activity statements within the fine-grained scheduling block on which to condition execution of the activity statement following the fine-grained scheduling block. The specific activity statements shall be selected randomly. Execution of the activity statement following the fine-grained scheduling block is scheduled after the selected activity statements.
   1) The *expression* shall be of an integer type. The value of the *expression* shall be determinable at solve time. If the value is 0, the **join_select** is equivalent to **join_none**.
   2) When an action array is traversed, each element of the array is considered a separate action that may be selected independently.

c) **join_none** specifies that the activity statement following the fine-grained scheduling block has no scheduling dependency on activity statements within the block.

d) **join_first** specifies that the activity statement following the fine-grained scheduling block has a runtime execution dependency on the first N activity statements within the fine-grained scheduling block to complete execution. The activity statement following the fine-grained scheduling block has no scheduling dependency on activity statements within the block, only a runtime dependency.
   1) The *expression* shall be of an integer type. The value of the *expression* shall be determinable at solve time. If the value is 0, the **join_first** is equivalent to **join_none**.
   2) When an action array is traversed, each element of the array is considered a separate action that may be selected independently.

The application scope of a fine-grained scheduling block is bounded by the sequential block that contains it. In other words, all activity statements that start within the fine-grained scheduling block shall complete before the statement following the containing sequential block begins. Activities started, but not joined,

Copyright © 2026 Accellera. All rights reserved.

within a fine-grained scheduling block are not implicitly waited for by any containing parallel or schedule blocks. Only the containing sequential block causes a join on activities started within it.

11.3.6.2 Examples

In Example 94, the innermost parallel block (`L4`) starts two activities (`L5` and `L6`), while only waiting for one (`L5`) to complete before continuing. Since `L5` traverses the action array `b`, all elements of `b` shall complete before continuing. The next level of parallel block (`L2`) waits for its two branches to complete (`L3` and `L4`), but does not wait for `L6` to complete. The outermost parallel block (`L1`) waits for one of its branches (`L2`) to complete before proceeding. This means that both `L7` and `L6` may be in-flight when `L8` is traversed.

```
B b[2];
activity {
   L1: parallel join_branch(L2) {
      L2: parallel {
         L3: do A;
         L4: parallel join_branch (L5) {
            L5: b;
            L6: do C;
         }
      }
      L7: do D;
   }
   L8: do F;
}
```

Example 94—Join_branch

The scheduling graph of the activity is shown in Figure 9.

Copyright © 2026 Accellera. All rights reserved.

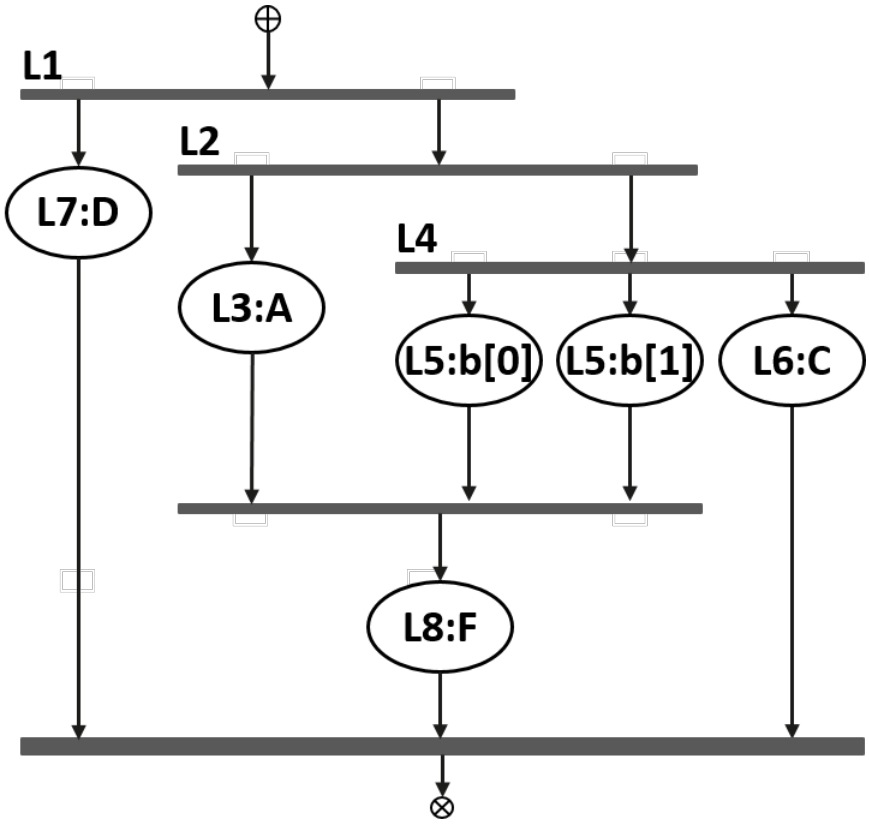


Figure 9—join_branch scheduling graph

The runtime behavior is shown in Figure 10.

Copyright © 2026 Accellera. All rights reserved.

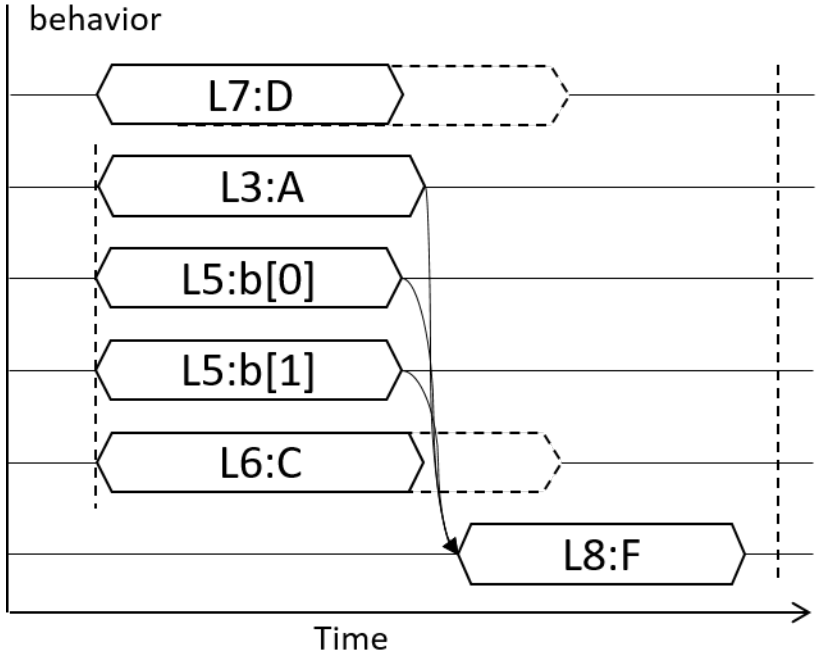


Figure 10—Join_branch runtime behavior

Activity scheduling in the presence of a fine-grained scheduling block is still subject to all other scheduling rules. For example, if both `L6` and `L8` in the example above contend for the same single resource, they shall be scheduled sequentially in order to avoid a resource conflict.

For the following four examples, assume that each of the three actions in the activity locks a resource from the same pool.

In Example 95, the **parallel** block causes traversal of branches `L1` and `L2` to be scheduled in parallel. The **join_branch** specifier causes traversal of action `C` to be scheduled with a sequential dependency on the activity statement labeled `L2`. Traversal of action `C` may not begin until the activity statement labeled `L2` has completed. To avoid adding additional scheduling dependencies, the resource pool would need a minimum of two resource instances. Actions `A` and `B` would each lock a resource instance, and `C`, since it is guaranteed not to start until `A` completes, would lock the same resource instance as that assigned to `A`. Note that this allocation is handled at solve-time, and is independent of whether `B` completes before or after `A` completes.

```
activity {
  L1 : parallel join_branch(L2) {
    L2: do A;
    L3: do B;
  }
  L4: do C;
}
```

Example 95—Join_branch with scheduling dependency

Copyright © 2026 Accellera. All rights reserved.

The scheduling graph of the activity is shown in Figure 11.

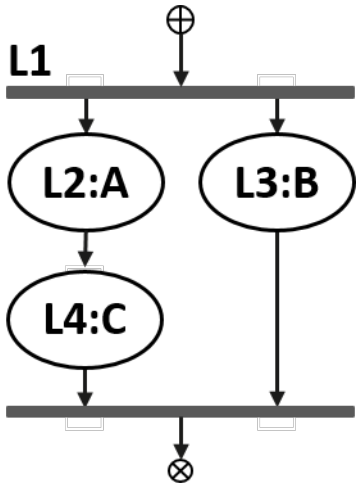


Figure 11—Scheduling graph of join_branch with scheduling dependency

The runtime behavior is shown in Figure 12.

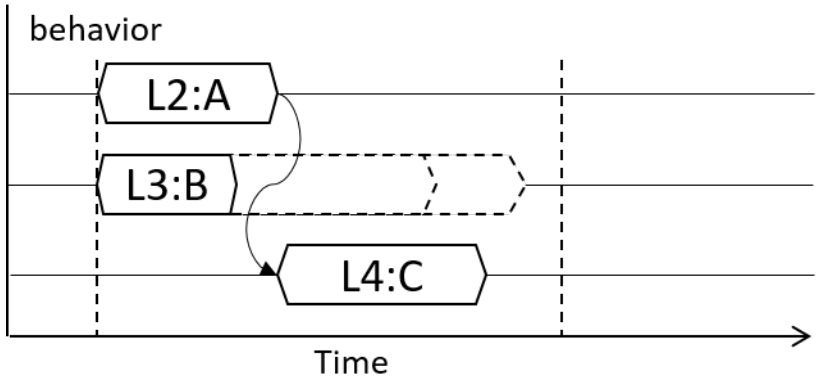


Figure 12—Runtime behavior of join_branch with scheduling dependency

In Example 96, the **parallel** block causes traversal of the branches labeled `L2` and `L3` to be scheduled in parallel. The **join_select** specifier causes traversal of action `C` to be scheduled with a sequential dependency on a random selection of either the branch labeled `L2` or `L3`. This means that traversal of `C` may not begin until after the selected target activity statement has completed. The tool randomly selects N (in this case, 1) target branch(es) from the candidate branches on which to make traversal of the following activity statement dependent.

In this example, the resource pool would need a minimum of two resource instances. Because the tool may not know which of `A` or `B` will complete first, it shall choose one and assign the same resource instance to action `C`. If the tool selected `L2` as the branch on which `C` depends, the behavior would be identical to the previous example.

Copyright © 2026 Accellera. All rights reserved.

```
activity {
  L1 : parallel join_select(1) {
    L2: do A;
    L3: do B;
  }
  L4: do C;
}
```

Example 96—Join_select

In Example 97, the **join_none** specifier causes traversal of action C to be scheduled with no dependencies. To avoid additional scheduling dependencies, the minimum size of the resource pool shall be three, since each action traversed in the activity shall have a unique resource instance.

Actions A and B are scheduled in parallel, and action C is scheduled concurrently with both of them. This means that C *could* start at the same time as A and B, but it may not. While the **parallel** statement precludes any dependencies between A and B, the **join_none** qualifier allows action C to be scheduled concurrently, but there may be additional dependencies between action C and action A and/or B.

```
activity {
  L1 : parallel join_none {
    L2: do A;
    L3: do B;
  }
  L4: do C;
}
```

Example 97—Join_none

The scheduling graph of the activity is shown in Figure 13.

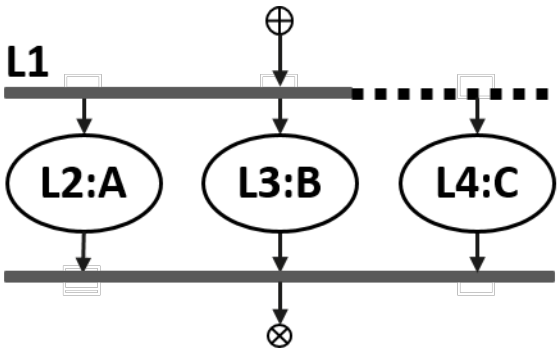


Figure 13—Join_none scheduling graph

In Example 98, the **join_first** specifier causes the PSS processing tool to condition execution of action C on runtime execution completion of the first of either action A or B. Since the scheduling tool may not know which action will complete first, there shall be a minimum of three resource instances in the pool in order to guarantee that C may execute immediately after whichever of A or B completes first. If there are two instances in the pool, the tool may assign either resource instance to C at solve-time. If the other action

Copyright © 2026 Accellera. All rights reserved.

assigned the same resource instance completes last, then action `C`, because it starts execution after the previous action completes, will also start its execution after the completion of the first action.

```
activity {
  L1 : parallel join_first(1) {
    L2: do A;
    L3: do B;
  }
  L4: do C;
}
```

Example 98—Join_first

The runtime behavior is shown in Figure 14.

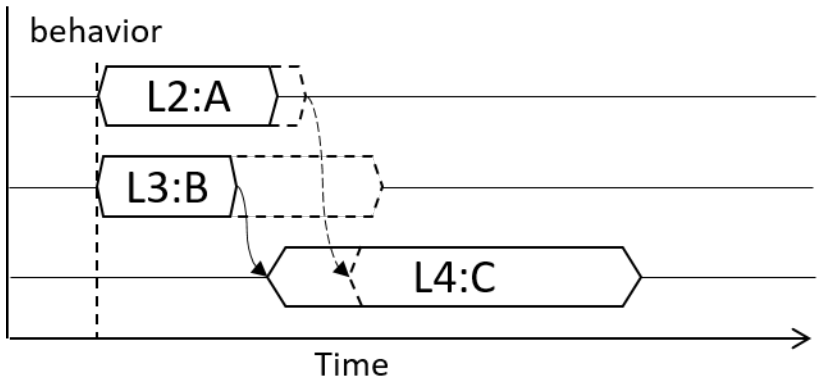


Figure 14—Runtime behavior of join_first

Example 99 illustrates how a **sequence** block bounds the impact of the fine-grained scheduling specifier. The execution of `L5` is scheduled in sequence with `L3`. `L4` and `L5` may be scheduled concurrently. `L6` is scheduled strictly sequentially to all statements inside `L1`, the **sequence** block.

```
activity {
   L1: sequence {
      L2: parallel join_branch(L3) {
         L3: do A;
         L4: do B;
      }
      L5: do C;
   }
   L6: do D;
}
```

Example 99—Scope of join inside sequence block

Copyright © 2026 Accellera. All rights reserved.

The scheduling graph is shown in Example 15.

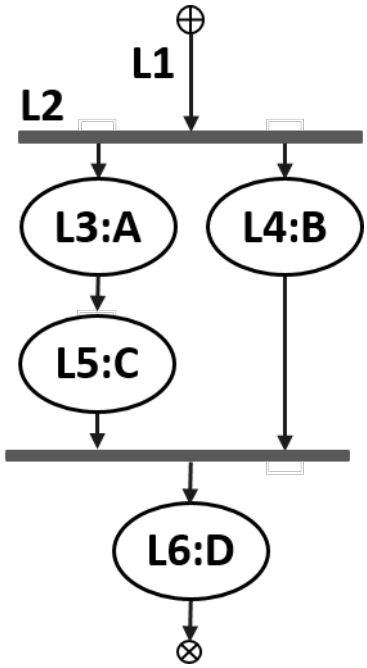


Figure 15—Scheduling graph of join inside sequence block

The runtime behavior is shown in Figure 16.

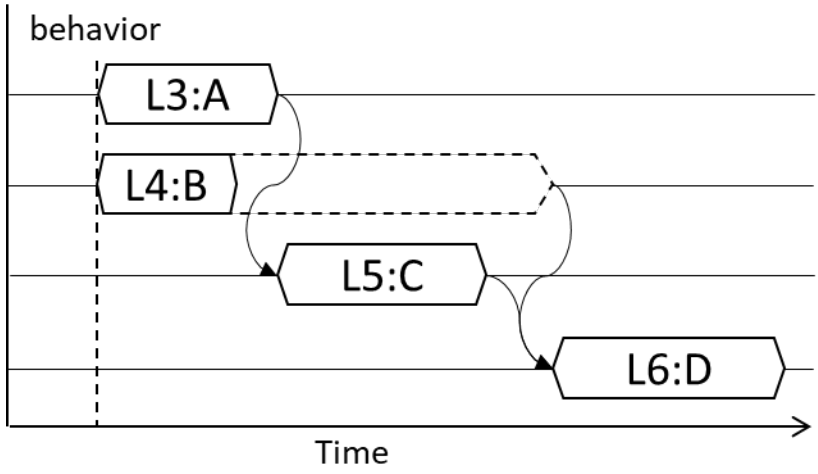


Figure 16—Runtime behavior of join inside sequence block

Copyright © 2026 Accellera. All rights reserved.

Example 100 shows how the **join** specification may also be used with the **schedule** block.

```
activity {
  L1 : schedule join_branch(L2) {
    L2: do A;
    L3: do B;
  }
  L4: do C;
}
```

Example 100—Specification for join with schedule block

Assuming there are no scheduling dependencies between actions `A` and `B`, the scheduling graph of **schedule** block `L1` is shown in Figure 17.

In all cases, action `C` is scheduled subsequent to action `A`. If `A` is scheduled before `B`, then `B` and `C` may–or may not–be scheduled concurrently, although there may be additional dependencies between them. If `B` is scheduled before `A`, the actions are executed in the order `B`, `A`, `C`. If `A` and `B` are scheduled concurrently, then `C` is still scheduled after `A`, but again may be concurrent with `B`, subject to any dependencies between `B` and `C`.

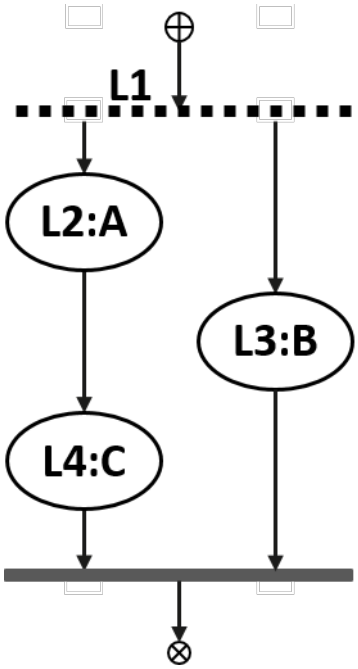


Figure 17—Scheduling graph join with schedule block

11.3.7 Atomic block specifier

Within an activity block, the *atomic block specifier* is used to preserve intended scheduling structure of its sub-activity, by preventing potential interference from other actions in the larger scenario. Example 101 and Example 102 in 11.3.7.2 demonstrate two typical causes for such interference: action inference and scheduling issues due to resource allocation. The atomic block specifier restricts the legal solution space by ruling out “unintended” (but otherwise legal) scheduling dependencies between actions within an atomic

Copyright © 2026 Accellera. All rights reserved.

block and the rest of the scenario. The following section defines which scheduling dependencies are ruled out and which remain legal.

An atomic block is analogous to an atomic action from a scheduling point of view, meaning that it can be substituted by an atomic action without change to the outside scheduling relations. All actions explicitly traversed in an atomic block are part of a single scheduling "cluster" (a nested subgraph of the scheduling dependency graph). In a transitive-reduced scheduling graph, the atomic block would have exactly one incoming edge and one outgoing edge. The incoming edge would represent "upward" dependencies, scheduling dependencies of an action traversed in the atomic block on outside actions. These outside actions become scheduling dependencies of the block as a whole (i.e., of all other actions in the cluster). The outgoing edge would represent "downward" dependencies, scheduling dependencies of an action within the cluster to an action outside the cluster. The outside action has a scheduling dependency on the block as a whole (i.e., on all other actions within the cluster).

11.3.7.1 Syntax

```
activity_atomic_block_stmt ::= atomic { { activity_stmt } }
```

Syntax 45—Atomic block

An *atomic set* is the set of all action executions corresponding to action traversal statements under the scope of an atomic block.

- This recursively includes all sub-actions of a compound action traversed in the atomic block.
- One atomic set can be a subset of another, but two atomic sets cannot have a non-empty intersection unless one is a subset of the other (this is guaranteed by the structure of activities).
- Inferred actions are never within an atomic set.

The following applies:

- If AS is an atomic set, $a_1 \in AS$, and $a_2 \notin AS$, then:
  1) If $a_1 \rightarrow a_2$, then for every $a_3 \in AS$, $a_3 \rightarrow a_2$; that is, if an action outside the atomic set has a scheduling dependency on an action inside the atomic set, then the outside action has a schedule dependency on all actions in the atomic set.
  2) If $a_2 \rightarrow a_1$ then for every $a_3 \in AS$, $a_2 \rightarrow a_3$; that is, if an action inside the atomic set has a scheduling dependency on an action outside the atomic set, then all actions in the atomic set have a scheduling dependency on the outside action.

11.3.7.2 Examples

Consider the code in Example 101. It demonstrates how the atomic specifier prevents the PSS solver from generating an unintended scenario scheduling due to the action inference process.

The atomic block specifier is used to ensure that `B_a` starts immediately after `A_a` completes. `B_a` may only start after `configX_a` completes. `configX_a` could require a meaningful amount of time to complete. `configX_a` needs to be inferred. Without the atomic specifier, `configX_a` could be inferred to execute after `A_a` and before `B_a`. With the atomic specifier, we are guaranteed a stress scenario where `B_a` is executed immediately after `A_a` completes.

Copyright © 2026 Accellera. All rights reserved.

```
action bringup_a {}

state config_s {
   rand mode_e mode;
}

action configX_a {
  output config_s out_cfg;
  constraint out_cfg.mode == X;
}

action A_a {}

action B_a {
  input config_s cfg;
  constraint cfg.mode == X;
}

action my_stress_seq_a {
  activity {
    do bringup_a;
    atomic {
      do A_a;
      do B_a;
    }
  }
}
```

Example 101—Atomic block to avoid action interference

Figure 18 illustrates undesired scheduling of the `configX_a` action when inferred, which can occur if the atomic specifier is not used.

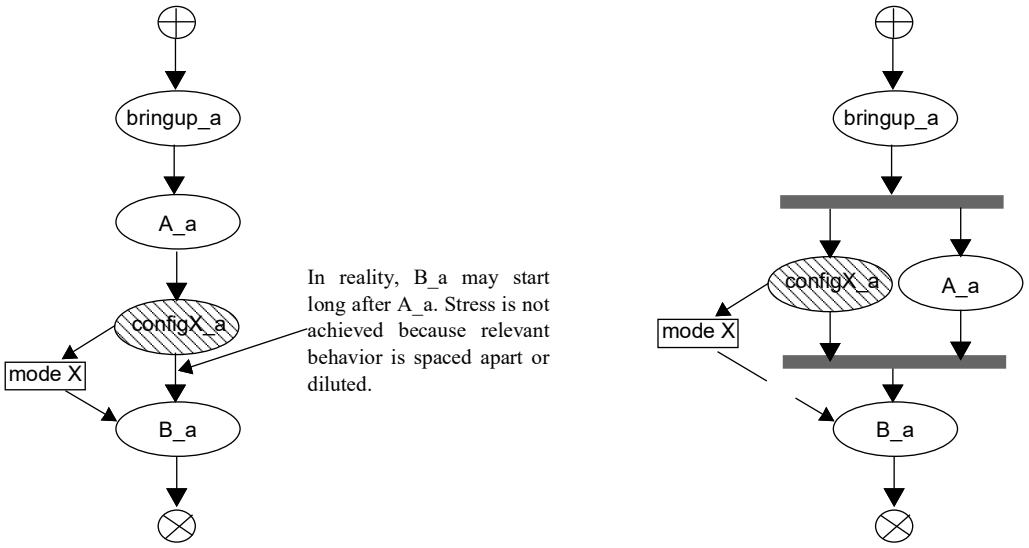


Figure 18—Scheduling graph of action interference

Copyright © 2026 Accellera. All rights reserved.

Figure 19 illustrates the cluster of actions in an atomic block (i.e., `A_a` and `B_a`) and how the `configX_a` action is an "upward" scheduling dependency of the atomic block. The figure shows two examples where `configX_a` is scheduled: a) after the bring-up and before the atomic block; b) in parallel with the bring-up and before the atomic block.

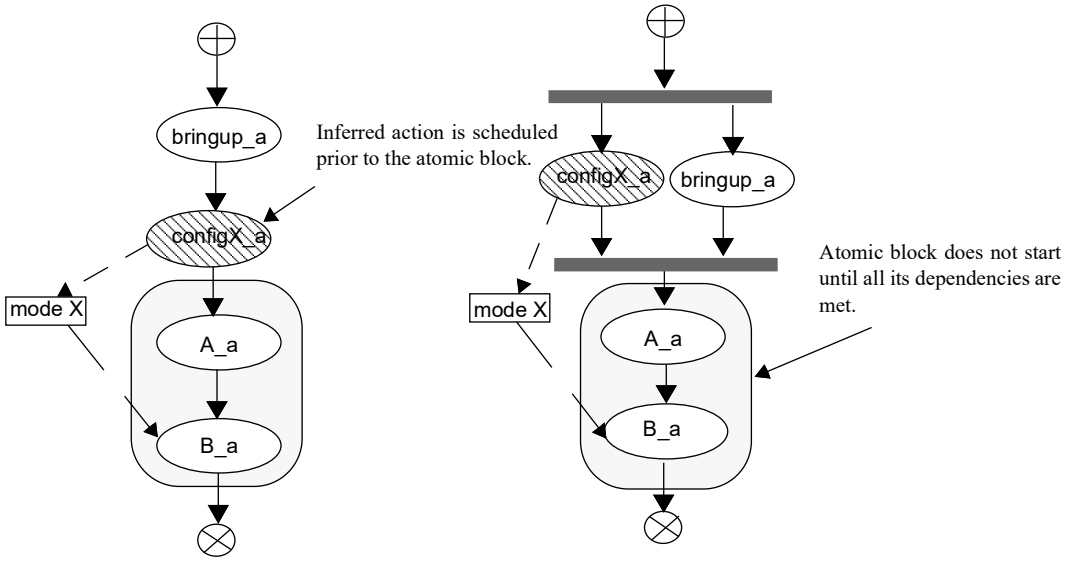


Figure 19—Scheduling graph of atomic block avoiding interference

Consider the code in Example 102. It demonstrates how the atomic specifier prevents the PSS solver from generating an unintended scenario scheduling due to a possible outcome of the resource allocation process.

Test intent of `my_stress_seq` is that `B` follows `A` as soon as possible. Figure 20 shows a scheduling solution that would violate this intent within the `my_test` scenario. `C` could be scheduled in parallel with `A` when both `B` and `C` happen to be assigned same resource slot, causing `B` to wait for completion of `C` which may take longer than `A`.

Copyright © 2026 Accellera. All rights reserved.

```
resource core_r {}
pool [4] core_r core_pool;

action A {}

action B {
  lock core_r core;
}

action C {
  lock core_r core;
}

action my_stress_seq {
  activity {
    atomic {
      do A;
      do B;
    }
  }
}

action my_test {
  activity {
    schedule {
      do my_stress_seq;
      do C;
    }
  }
}
```

Example 102—Atomic block to avoid resource allocation issues

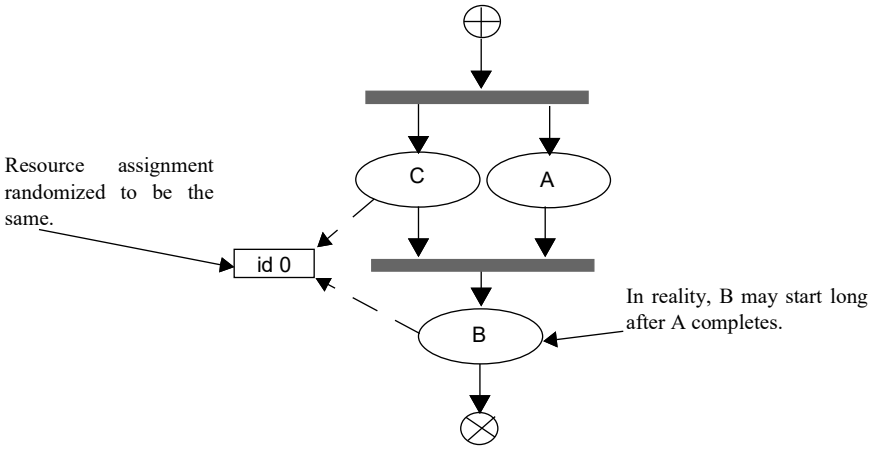


Figure 20—Scheduling graph of resource allocation issues

Copyright © 2026 Accellera. All rights reserved.

11.4 Activity control flow constructs

In addition to defining sequential and parallel blocks of action execution, repetition and branching statements can be used inside the **activity** clause.

11.4.1 Repeat (count)

The **repeat** statement allows the specification of a loop consisting of one or more actions inside an activity. This section describes the *count-expression* variant (see Syntax 46) and 11.4.2 describes the *while-expression* variant.

11.4.1.1 Syntax

```
activity_repeat_stmt ::=
    repeat ( [ index_identifier : ] expression ) activity_stmt
  | . . .
```

Syntax 46—Repeat-count statement

The following also apply:

a) *expression* shall be a non-negative integer expression (**int** or **bit**).

b) Intuitively, the *activity_stmt* is iterated the number of times specified in the *expression*. An optional index-variable identifier can be specified that ranges between `0` and one less than the iteration count. If the expression evaluates to `0`, the *activity_stmt* is not evaluated at all.

c) Formally, the **repeat-**count statement specifies sequential scheduling between *N* sets of action executions per the evaluation of *activity_stmt N* times, where *N* is the number to which *expression* evaluates (see 6.3.2).

d) The choice of values to **rand** attributes figuring in the *expression* shall be such that it yields legal execution scheduling.

11.4.1.2 Examples

In Example 103, the resulting execution is six sequential action executions, alternating `A`'s and `B`'s, with five scheduling dependencies: $(A_0) \rightarrow (B_0)$, $(B_0) \rightarrow (A_1)$, $(A_1) \rightarrow (B_1)$, $(B_1) \rightarrow (A_2)$, $(A_2) \rightarrow (B_2)$.

```
action  my_test {
   A a;
   B b;
   activity {
      repeat (3) {
         a;
         b;
      }
   }
};
```

Example 103—Repeat statement

Copyright © 2026 Accellera. All rights reserved.

Example 104 shows an additional example of using **repeat**-count.

```
action my_test {
    my_action1 action1;
    my_action2 action2;
    activity {
      repeat (i : 10) {
         if ((i % 4) == 0) {
           action1;
         } else {
           action2;
         }
      }
   }
};
```

Example 104—Another repeat statement

11.4.2 Repeat-while

The **repeat** statement allows the specification of a loop consisting of one or more actions inside an activity. This section describes the *while-expression* variant (see Syntax 47).

11.4.2.1 Syntax

```
activity_repeat_stmt ::=
  . . .
  | repeat activity_stmt while ( expression ) ;
```

Syntax 47—Repeat-while statement

The following also apply:

a) *expression* shall be of type **bool**.

b) Intuitively, the *activity_stmt* is iterated so long as the *expression* condition is *true*, as sampled after the *activity_stmt*.

c) Formally, the **repeat-while** statement specifies sequential scheduling between multiple sets of action executions per the iterative evaluation of *activity_stmt*. The evaluation of *activity_stmt* continues repeatedly so long as *expression* evaluates to *true*. *expression* is evaluated after the execution of each set in the **repeat-while** block.

Copyright © 2026 Accellera. All rights reserved.

11.4.2.2 Example

```
component top {

    function bit is_last_one();

    action do_something {
        bit last_one;

        exec post_solve {
            last_one = comp.is_last_one();
        }

        exec body C = """
            printf("Do Something\n");
        """;
    }

    action entry {
        do_something s1;

        activity {
            repeat {
                s1;
            } while (s1.last_one !=0);
        }
    }
}
```

Example 105—Repeat-while statement

11.4.3 Foreach

The **foreach** construct iterates over the elements of a collection (see Syntax 48). See also Example 106.

11.4.3.1 Syntax

activity_foreach_stmt ::=
foreach ([*iterator*_identifier **:**] expression [**[** *index*_identifier **]**] **)** activity_stmt

Syntax 48—Foreach statement

The following also apply:

a) *expression* shall be of a collection type (i.e., **array**, **list**, **map** or **set**), including fixed-sized arrays of *modeling elements* (not including monitors).

b) The body of the **foreach** statement is a sequential block in which *activity_stmt* is evaluated once for each element in the collection.

c) *iterator_identifier* specifies the name of an iterator variable of the collection element type. Within *activity_stmt*, the iterator variable, when specified, is an alias to the collection element of the current iteration.

d) *index_identifier* specifies the name of an index variable. Within *activity_stmt,* the index variable, when specified, corresponds to the element index of the current iteration.

Copyright © 2026 Accellera. All rights reserved.

1) For **array**s and **list**s, the index variable shall be a variable of type **int**, ranging from **0** to one less than the size of the collection variable, in that order.
2) For **map**s, the index variable shall be a variable of the same type as the **map** keys, and range over the values of the keys. The order of key traversal is undetermined.
3) For **set**s, an index variable shall not be specified.

e) Both the index and iterator variables, if specified, are implicitly declared within the **foreach** scope and limited to that scope. Regular name resolution rules apply when the implicitly declared variables are used within the **foreach** body. For example, if there is a variable in an outer scope with the same name as the index variable, that variable is shadowed (masked) by the index variable within the **foreach** body. The index and iterator variables are not visible outside the **foreach** scope.

f) Either an index variable or an iterator variable or both shall be specified. For a **set**, an iterator variable shall be specified, but not an index variable.

11.4.3.2 Example

```
action my_action1 {
    rand bit[4] val;
    // ...
}

action my_test {
    rand bit[4] in [0..7] a[16];
    my_action1 action1;

    activity {
      foreach (a[j]) {
        action1 with {val <= a[j];};
      }
    }
};
```

Example 106—Foreach statement

11.4.4 Select

The **select** statement specifies a branch point in the traversal of the activity (see Syntax 49).

11.4.4.1 Syntax

```
activity_select_stmt ::= select { select_branch select_branch { select_branch } }
select_branch ::= [ [ ( expression ) ] [ [ expression ] ] : ] activity_stmt
```

Syntax 49—Select statement

The following also apply:

a) Intuitively, a **select** statement executes one out of a number of possible activities.

b) One or more of the *activity_stmt*s may optionally have a guard condition specified in parentheses (( )). Guard condition expressions shall be of Boolean type. When the **select** statement is evaluated, only those *activity_stmt*s whose guard condition evaluates to *true* or that do not have a guard condition are considered enabled.

Copyright © 2026 Accellera. All rights reserved.

c) Formally, each evaluation of a **select** statement corresponds to the evaluation of just one of the *select_branch* statements. All scheduling requirements shall hold for the selected **activity** statement.

d) Optionally, all *activity_stmt*s may include a *weight expression*, which is a numeric expression that evaluates to a non-negative integer value. The probability of choosing an enabled *activity_stmt* is the weight of the given statement divided by the sum of the weights of all enabled statements. If the *activity_stmt* is an array of action handles, then the *weight expression* is assigned to each element of the array, from which one element is selected and traversed.

e) If any *activity_stmt* has a *weight expression*, then any statement without an explicit *weight expression* associated with it shall have a weight of `1`.

f) It shall be illegal if no **activity** statement is valid according to the active constraint and scheduling requirements and the evaluation of the guard conditions.

11.4.4.2 Examples

In Example 107, the **select** statement causes the activity to select `action1` or `action2` during each execution of the activity.

```
action my_test {
  my_action1          action1;
  my_action2          action2;
  activity {
    select {
      action1;
      action2;
    }
  }
}
```

Example 107—Select statement

In Example 108, the branch selected shall depend on the value of `a` when the **select** statement is evaluated.

a) `a==0` means that all three branches could be chosen, according to their weights.
   1) `action1` is chosen with a probability of 20%.
   2) `action2` is chosen with a probability of 30%.
   3) `action3` is chosen with a probability of 50%.

b) `a in [1..3]` means that `action2` or `action3` is traversed according to their weights.
   1) `action2` is chosen with a probability of 37.5%.
   2) `action3` is chosen with a probability of 62.5%.

c) `a==4` means that only `action3` is traversed.

Copyright © 2026 Accellera. All rights reserved.

```
action my_test {
  my_action1 action1;
  my_action2 action2;
  my_action3 action3;
  rand int in [0..4] a;
  activity {
    select {
      (a == 0     )[20]: action1;
      (a in [0..3])[30]: action2;
                   [50]: action3;
    }
  }
}
```

Example 108—Select statement with guard conditions and weights

In Example 109, the **select** statement causes the activity to select `action1` or one element of `action2` during the execution of the **activity**. Since the weight expression of `2` is applied to each element of the `action2` array, there is a 40% chance that either element of that array is chosen, and a 20% (weight of `1`) chance of choosing `action1`.

```
action my_test {
  my_action1 action1;
  my_action2 action2[2];

  activity {
    select {
           action1;
      [2]: action2;
    }
  }
}
```

Example 109—Select statement with array of action handles

11.4.5 If-else

The **if-else** statement introduces a branch point in the traversal of the activity (see Syntax 50).

11.4.5.1 Syntax

activity_if_else_stmt ::= **if (** expression **)** activity_stmt [**else** activity_stmt]

Syntax 50—If-else statement

The following also apply:

a) *expression* shall be of type **bool**.

b) Intuitively, an **if-else** statement executes some activity if a condition holds, and, otherwise (if specified), the alternative activity.

c) Formally, the **if-else** statement specifies the scheduling of the set of action executions per the evaluation of the first *activity_stmt* if *expression* evaluates to *true* or the second *activity_stmt* (following **else**) if present and *expression* evaluates to *false*.

Copyright © 2026 Accellera. All rights reserved.

d) The scheduling relationships need only be met for one branch for each evaluation of the activity.

e) The choice of values to **rand** attributes figuring in the *expression* shall be such that it yields legal execution scheduling.

11.4.5.2 Example

If the scheduling requirements for Example 110 required selection of the `b` branch, then the value selected for `x` shall be `<= 5`.

```
action my_test {
   rand int in [1..10] x;
   A a;
   B b;
   activity {
      if (x > 5)
         a;
      else
         b;
   }
};
```

Example 110—If-else statement

11.4.6 Match

The **match** statement specifies a multi-way decision point in the traversal of the activity that tests whether an expression matches any of a number of other expressions and traverses one of the matching branches accordingly (see Syntax 51).

11.4.6.1 Syntax

```
activity_match_stmt ::= match ( match_expression ) { match_choice { match_choice } }
match_expression ::= expression
match_choice ::=
  [ open_range_list ] : activity_stmt
 | default : activity_stmt
```

Syntax 51—Match statement

The following also apply:

a) When the **match** statement is executed, the *match_expression* is evaluated.

b) After the *match_expression* is evaluated, the *open_range_list* of each *match_choice* shall be compared to the *match_expression*. *open_range_list*s are described in 8.5.9.1.

c) If there is exactly one match, then the corresponding branch shall be traversed.

d) If there is more than one match, then one of the matching *match_choice*s shall be randomly traversed.

e) If there are no matches, then the **default** branch, if provided, shall be traversed.

f) The **default** branch is optional. There may be at most one **default** branch in the **match** statement.

Copyright © 2026 Accellera. All rights reserved.

g) As with a **select** statement, it shall be an error if no *match_choice* is valid according to the active constraint and scheduling requirements and the evaluation of the *match_expression* against the *match_choice open_range_list*s.

11.4.6.2 Example

In Example 111, the **match** statement causes the **activity** to evaluate the data field `in_security_data.val` and select a branch according to its value at each execution of the activity. If the data field is equal to `LEVEL2`, `action1` is traversed. If the data field is equal to `LEVEL5`, `action2` is traversed. If the data field is equal to `LEVEL3` or `LEVEL4`, then either `action1` or `action2` is traversed at random. For any other value of the data field, `action3` is traversed.

```
action my_test {
  rand security_data in_security_data;
  my_action1 action1;
  my_action2 action2;
  my_action3 action3;
  activity {
    match (in_security_data.val) {
      [LEVEL2..LEVEL4]:
        action1;
      [LEVEL3..LEVEL5]:
        action2;
      default:
        action3;
    }
}
```

Example 111—Match statement

11.5 Activity construction statements

11.5.1 Replicate

The **replicate** statement is a generative activity statement interpreted as an in-place expansion of a specified statement multiple times. The **replicate** statement does not introduce an additional layer of scheduling or control flow. The execution semantics applied to the expanded statements depend on the context. In particular, replicating a statement N times under a **parallel** statement executes the same statement N times in parallel. Unlike a **repeat** statement, **replicate** provides a way to reference specific expansion instances from above using a label array.

11.5.1.1 Syntax

```
activity_replicate_stmt ::=
  replicate ( [ index_identifier : ] expression ) [ label_identifier [ ] : ] labeled_activity_stmt
```

Syntax 52—Replicate statement

The following also apply:

a) *expression* shall be a positive integer expression (**int** or **bit**).

Copyright © 2026 Accellera. All rights reserved.

b) The **replicate** statement expands in-place to *labeled_activity_stmt* replicated the number of times specified in the *expression*. An optional index variable *index_identifier* may be specified that ranges between `0` and one less than the iteration count.

c) The execution semantics of a **replicate** statement where *expression* evaluates to N are equivalent to the execution semantics of N occurrences of *labeled_activity_stmt* directly traversed in its enclosing activity scope.

d) The number of replications shall be known as part of the solve process. In other words, *expression* may not contain an attribute that is assigned in the context of a runtime *exec block* (**body**/**run_start**/**run_end**).

e) A *label_identifier* may optionally be used to label the replicated statement in the form of a label array. If used, each expanded occurrence of *labeled_activity_stmt* becomes a named sub-activity with the label *label_identifier*[0] … *label_identifier*[N-1] respectively, where N is the number of expanded occurrences. Reference can be made to labels and action handles declared under the **replicate** and its nested scopes using array indexing on the label. (See more on hierarchical activity references in 11.8).

f) Labels may be used to name sub-activities inside the scope of a **replicate** statement only if the *label_identifier* is specified. A label under a **replicate** statement without a named label array leads to name conflict between the replicated sub-activities (see scoping rules for named sub-activities in 11.8.2).

g) Traversing a named action handle within a **replicate** scope that is declared outside the **replicate** scope shall not result in multiple traversal when the **replicate** statement is expanded (see 11.3.1.1(c)). Both anonymous action traversal and action traversal of an action handle declared locally inside the **replicate** scope are allowed.

11.5.1.2 Examples

In Example 112, the resulting execution is either two, three, or four parallel executions of the sequence `A -> B`.

```
action my_test {
  rand int in [2..4] count;
  activity {
    parallel {
      replicate (count) {
        do A;
        do B;
      }
    }
  }
};
```

Example 112—Replicate statement

Copyright © 2026 Accellera. All rights reserved.

In Example 113, the execution of action `my_test` results in one execution of `A` as well as four executions of `B`, all in the scope of the **schedule** statement, that is, invoked in any order that satisfies the scheduling rules.

```
action my_test {
  activity {
    schedule {
      do A;
      replicate (i: 4) do B with { size == i*10; };
    }
  }
};
```

Example 113—Replicate statement with index variable

Example 113 can be rewritten in the following equivalent way to eliminate the **replicate** statement:

```
action my_test {
  activity {
    schedule {
      do A;
      do B with { size == 0*10; };
      do B with { size == 1*10; };
      do B with { size == 2*10; };
      do B with { size == 3*10; };
    }
  }
};
```

Example 114—Rewriting previous example without replicate statement

Copyright © 2026 Accellera. All rights reserved.

Example 115 illustrates the use of a **replicate** label array for unique hierarchical paths to specific expansion instances. References are made to action handles declared and traversed in specific expansion instances of a **replicate** statement from outside its scope.

```
action my_compound {
  rand int in [2..4] count;
  activity {
    parallel {
      replicate (count) RL[]: {
        A a;
        B b;
        a;
        b;
      }
    }
    if (RL[count-1].a.x ==0) { // 'a' of the last replicate expansion
      do C;
    }
  }
};

action my_test {
  activity {
    do my_compound with {
      RL[0].a.x == 10; // 'a' of the first replicate expansion
    };
  }
};
```

Example 115—Replicate statement with label array

In Example 116 a number of error situations are demonstrated. Note that label `L` in this example causes a name conflict between the named sub-activities in the expansion of the **replicate** statement (see also 11.8.2).

Copyright © 2026 Accellera. All rights reserved.

```
action my_test {
  A a;
  C c_arr[4];
  activity {
    schedule {
      replicate (i:4) {
        B b;
        a;          // Error - traversal of action handle
                    // declared outside the replicate scope
        b;          // OK - action handle declared inside replicate scope
        c_arr[i];   // OK - each element of the action handle array is a
                    // unique action handle, so does not cause the same
                    // handle to be traversed multiple times
        L: select { // Error - label causes name conflict in expansion
          do A;
          do B;
        }
      }
    }
  }
};
```

Example 116—Replicate statement error situations

11.6 Activity evaluation with extension and inheritance

Compound actions support both type inheritance and type extension (see Clause 17). When type extension is used to contribute one or more activities to an action type, the execution semantics are the same as if all the contributed activities were scheduled along with all the activities from the initial definition.

In Example 117, action type `entry` traverses action type `A`. Extensions to action type `entry` include activities that traverse action types `B` and `C`.

Copyright © 2026 Accellera. All rights reserved.

```
component pss_top {
    action A { };
    action B { };
    action C { };

    action entry {
        activity {
            do A;
        }
    }

    extend action entry {
        activity {
            do B;
        }
    }

    extend action entry {
        activity {
            do C;
        }
    }
}
```

Example 117—Extended action traversal

The semantics of **activity** in the presence of type extension state that all three activity blocks will be traversed in an implied **schedule** block. In other words, Example 117 is equivalent to the hand-coded example shown in Example 118.

```
component pss_top {
    action A { };
    action B { };
    action C { };

    action entry {
        activity {
            schedule {
                do A;
                do B;
                do C;
            }
        }
    }
}
```

Example 118—Hand-coded action traversal

When a compound action inherits from another compound action, any activities declared in the inheriting action *shadow* (mask) the activity (or activities) declared in the base action. The "**super;**" statement can be used to traverse the activity (or activities) declared in the base action. The "**super;**" activity statement implicitly defines a scope corresponding to the parent activity: a sequence scope for each activity block and, when multiple activity blocks are present, an additional schedule scope. The "**super;**" statement shall not be labeled.

Copyright © 2026 Accellera. All rights reserved.

Activity paths declared in the base action may be referenced using the “super” prefix. The “super” prefix may be omitted when the reference is unambiguous, analogous to access rules for fields inherited from a base action.

In Example 119, action `base` declares an activity that traverses action type `A`. Action `ext1` inherits from `base` and replaces the activity declared in `base` with an activity that traverses action type `B`. Action `ext2` inherits from `base` and replaces the activity declared in `base` with an activity that first traverses the activity declared in `base`, then traverses action type `C`.

```
component pss_top {
    action A { }
    action B { }
    action C { }

    action base {
        activity {
            do A;
        }
    }

    action ext1 : base {
        activity {
            do B;
        }
    }

    action ext2 : base {
        activity {
            super;
            do C;
        }
    }
}
```

Example 119—Inheritance and traversal

11.7 Symbols

To assist in reuse and simplify the specification of repetitive behaviors in a single activity, a *symbol* may be declared to represent a subset of activity functionality (see Syntax 53). The **symbol** may be used as a node in the activity.

A **symbol** may activate another **symbol**, but **symbol**s are not recursive and may not activate themselves.

11.7.1 Syntax

```
symbol_declaration ::= symbol symbol_identifier [ ( symbol_paramlist ) ] { { activity_stmt } }
symbol_paramlist ::= [ symbol_param { , symbol_param } ]
symbol_param ::= data_type identifier
```

Syntax 53—Symbol declaration

Copyright © 2026 Accellera. All rights reserved.

11.7.2 Examples

Example 120 depicts using a symbol. In this case, the desired activity is a sequence of choices between `a`*N* and `b`*N*, followed by a sequence of `c`*N* actions. This statement could be specified in-line, but for brevity of the top-level activity description, a symbol is declared for the sequence of `a`*N* and `b`*N* selections. The symbol is then referenced in the top-level activity, which has the same effect as specifying the `a`*N*/`b`*N* sequence of selects in-line.

```
component entity {
   action a { }
   action b { }
   action c { }

   action top {
      a a1, a2, a3;
      b b1, b2, b3;
      c c1, c2, c3;

      symbol a_or_b {
         select {a1; b1; }
         select {a2; b2; }
         select {a3; b3; }
      }

      activity {
         a_or_b;
         c1;
         c2;
         c3;
      }
   }
}
```

Example 120—Using a symbol

Copyright © 2026 Accellera. All rights reserved.

Example 121 depicts using a parameterized symbol.

```
component entity {
  action a { }
  action b { }
  action c { }
  action top {
    a a1, a2, a3;
    b b1, b2, b3;
    c c1, c2, c3;
    symbol ab_or_ba (a aa, b bb) {
      select {
        { aa; bb; }
        { bb; aa; }
      }
    }
    activity {
      ab_or_ba(a1,b1);
      ab_or_ba(a2,b2);
      ab_or_ba(a3,b3);
      c1;
      c2;
      c3;
    }
  }
}
```

Example 121—Using a parameterized symbol

11.8 Named sub-activities

Sub-activities are structured elements of an activity. Naming sub-activities is a way to specify a logical tree structure of sub-activities within an activity. This tree serves for making hierarchical references, both to action-handle variables declared in-line, as well as to the **activity** statements themselves. The hierarchical paths thus exposed abstract from the concrete syntactic structure of the activity, since only explicitly labeled statements constitute a new hierarchy level.

11.8.1 Syntax

A named sub-activity is declared by labeling an **activity** statement, see Syntax 39.

11.8.2 Scoping rules for named sub-activities

Activity statement labels shall be unique in the context of the containing named sub-activity—the nearest lexically-containing statement which is labeled. Activity statement labels shall not conflict with local variable names, including named action handles. Unlabeled activity statements do not constitute a separate naming scope for sub-activities.

Note that labeling activity statements inside the scope of a **replicate** statement leads to name conflicts between the expanded sub-activities, unless a label array is specified (see 11.5.1.1). With a **replicate** label array, each expanded named sub-activity has a unique hierarchical path.

In Example 122, some **activity** statements are labeled while others are not. The second occurrence of label `L2` is conflicting with the first because the **if** statement under which the first occurs is not labeled and hence is not a separate naming scope for sub-activities.

Copyright © 2026 Accellera. All rights reserved.

```
action A {};

action B {
  int x;
  activity {
    L1: parallel { // 'L1' is 1st level named sub-activity
      if (x > 10) {
        L2: {      // 'L2' is 2nd level named sub-activity
          A a;
          a;
        }
        {
          A a; // OK - this is a separate naming scope for variables
          a;
        }
      }
      L2: { // Error - this 'L2' conflicts with 'L2' above
        A a;
        a;
      }
    }
  }
};
```

Example 122—Scoping and named sub-activities

Example 123 below demonstrates a name conflict between a local action-handle variable and a label of an activity statement in the same named sub-activity. This is not allowed, as it would render the hierarchical path `L.a` from **action** `A`'s scope ambiguous.

```
action A {
  activity {
    L: schedule {
         A a;
         B b;
         a: { // illegal label!
           do C;
           do D;
         }
       }
  }
  constraint parallel {L.a, L.b};
}
```

Example 123—Label name conflict for activity statement

11.8.3 Hierarchical references using named sub-activity

Named sub-activities, introduced through labels, allow referencing action-handle variables using hierarchical paths. References can be made to a variable from within the same activity, from the compound action top-level scope, and from outside the action scope.

Copyright © 2026 Accellera. All rights reserved.

A hierarchical activity path uses labels in a way similar to variables of struct and array types. The dot operator (.) in the case of simple labels, or the indexing operator (**[]**) and other array operators in the case of label arrays (introduced by **replicate** statements), may be used to reference named sub-activity blocks.

Only action handles declared directly under a labeled activity statement can be accessed outside their direct lexical scope. Action handles declared in an unnamed activity scope cannot be accessed from outside that scope.

Note that the top activity scope is unnamed. For an action handle to be directly accessible in the top-level action scope, or from outside the current scope, it shall be declared at the top-level action scope.

In Example 124, **action** `B` declares action-handle variables in labeled activity statement scopes, thus making them accessible from outside by using hierarchical paths. **action** `C` uses hierarchical paths to constrain the sub-actions of its sub-actions `b1` and `b2`.

```
action A { rand int x; };

action B {
  A a;
  activity {
    a;
    my_seq: sequence {
      A a;
      a;
      parallel {
        my_rep: repeat (3) {
          A a;
          a;
        };
        sequence {
          A a; // this 'a' is declared in unnamed scope
          a;   // can't be accessed from outside
        };
      };
    };
  };
};

action C {
  B b1, b2;
  constraint b1.a.x == 1;
  constraint b1.my_seq.a.x == 2;
  constraint b1.my_seq.my_rep.a.x == 3; // applies to all three iterations
                                        // of the loop
  activity {
    b1;
    b2 with { my_seq.my_rep.a.x == 4; }; // likewise
  }
};
```

Example 124—Hierarchical references and named sub-activities

Copyright © 2026 Accellera. All rights reserved.

11.9 Explicitly binding flow objects

Input and output fields of **action**s may be explicitly connected to actions using the **bind** statement (see Syntax 54). It states that the fields of the respective **action**s reference the same object—the output of one action is the input of another.

11.9.1 Syntax

```
activity_bind_stmt ::= bind hierarchical_id activity_bind_item_or_list ;
activity_bind_item_or_list ::=
    hierarchical_id
  | { hierarchical_id_list }
```

Syntax 54—Bind statement

The following also apply:

a) Reference fields that are bound shall be of the same object type.

b) Explicit binding shall conform to the scheduling and connectivity rules of the respective flow object kind defined in 9.3.4.

c) Explicit binding can only associate reference fields that are statically bound to the same pool instance (see 12.3).

d) The order in which the fields are listed does not matter.

Copyright © 2026 Accellera. All rights reserved.

11.9.2 Examples

Examples of binding are shown in Example 125.

```
component top{
    buffer B {rand int a;};
    action P1 {
        output B out;
    };
    action P2 {
        output B out;
    };
    action C {
        input B inp;
    };

    pool B B_p;
    bind B_p {*};

    action T {
        P1 p1;
        P2 p2;
        C c;
        activity {
            p1;
            p2;
            c;
            bind p1.out c.inp; // c.inp.a == p1.out.a
        };
    }
};
```

Example 125—Bind statement

11.10 Hierarchical flow object binding

As discussed in 9.3.4, actions, including compound actions, may declare inputs and/or outputs of a given flow object type. When a compound action has inputs and/or outputs of the same type and direction as its sub-action and which are statically bound to the same pool (see 12.3), the **bind** statement may be used to associate the compound action's input/output with the desired sub-action input/output. The compound action's input/output shall be the first argument to the **bind** statement.

The outermost compound action that declares the input/output determines its scheduling implications, even if it binds the input/output to that of a sub-action. The binding to a corresponding input/output of a sub-action simply delegates the object reference to the sub-action.

In the case of a buffer object input to the compound action, the action that produces the buffer object shall complete before the activity of the compound action begins, regardless of where within the activity the sub-action to which the input buffer is bound begins. Similarly, the compound action's activity shall complete before the compound action's output buffer is available, regardless of where in the compound action's activity the sub-action that produces the buffer object executes. The corollary to this statement is that no other sub-action in the compound action's activity may have an input explicitly hierarchically bound to the compound action's buffer output object. Similarly, no sub-action in the compound action's activity may have an output that is explicitly hierarchically bound to the compound action's input object. Consider Example 126.

Copyright © 2026 Accellera. All rights reserved.

```
action sub_a {
  input data_buf din;
  output data_buf dout;
}

action compound_a {
  input data_buf data_in;
  output data_buf data_out;
  sub_a a1, a2;
  activity {
    a1;
    a2;
    bind a1.dout a2.din;
    bind data_in a1.din;   // hierarchical bind
    bind data_out a2.dout; // hierarchical bind
//  The following bind statements would be illegal
//  bind data_in a1.dout;  // sub-action output may not be bound to
//                         // compound action's input
//  bind data_out a2.din;  // sub-action input may not be bound to
//                         // compound action's output
  }
}
```

Example 126—Hierarchical flow binding for buffer objects

For stream objects, the compound action's activity shall execute in parallel with the action that produces the input stream object to the compound action or consumes the stream object output by the compound action. A sub-action within the activity of a compound action that is bound to a stream input/output of the compound action shall be an initial action in the activity of the compound action. Consider Example 127.

```
action sub_a {
  input data_str din;
  output data_buf dout;
}

action compound_a {
  input data_str data_in;
  output data_buf data_out;
  sub_a a1, a2;
  activity {
    a1;
    a2;
    bind data_in a1.din;  // hierarchical bind
//  The following bind statement would be illegal
//  bind data_in a2.din;  // a2 is not scheduled in parallel with compound_a
  }
}
```

Example 127—Hierarchical flow binding for stream objects

For state object outputs of the compound action, the activity shall complete before any other action may write to or read from the state object, regardless of where in the activity the sub-action executes within the activity. Only one sub-action may be bound to the compound action's state object output. Any number of sub-actions may have input state objects bound to the compound action's state object input.

Copyright © 2026 Accellera. All rights reserved.

11.11 Hierarchical resource object binding

As discussed in 9.4.2, actions, including compound actions, may claim a resource object of a given type. When a compound action claims a resource of the same type as its sub-action(s) and where the compound action and the sub-action are bound to the same pool, the **bind** statement may be used to associate the compound action's resource with the desired sub-action resource. The compound action's resource shall be the first argument to the **bind** statement.

The outermost compound action that claims the resource determines its scheduling implications. The binding to a corresponding resource of a sub-action simply delegates the resource reference to the sub-action.

The compound action's claim on the resource determines the scheduling of the compound action relative to other actions and that claim is valid for the duration of the activity. The sub-actions' resource claim determines the relative scheduling of the sub-actions in the context of the activity. In the absence of the explicit resource binding, the compound action and its sub-action(s) claim resources from the pool to which they are bound. Thus, it shall be illegal for a sub-action to lock the same resource instance that is locked by the compound action.

A resource locked by the compound action may be bound to any resource(s) in the sub-action(s). Thus, only one sub-action that locks the resource reference may execute in the activity at any given time and no sharing sub-actions may execute at the same time. If the resource that is locked by the compound action is bound to a shared resource(s) in the sub-action(s), there is no further scheduling dependency.

A resource shared by the compound action may only be bound to a shared resource(s) in the sub-action(s). Since the compound action's shared resource may also be claimed by another action, there is no way to guarantee exclusive access to the resource by any sub-action; so, it shall be illegal to bind a shared resource to a locking sub-action resource.

In Example 128, the compound action locks resources `crlkA` and `crlkB`, so no other actions outside of `compound_a` may lock either resource for the duration of the activity.

```
action sub_a {
  lock  res_r rlkA, rlkB;
  share res_r rshA, rshB;
}

action compound_a {
  lock  res_r crlkA, crlkB;
  share res_r crshA, crshB;
  sub_a a1, a2;
  activity {
    schedule {
      a1;
      a2;
    }
    bind crlkA {a1.rlkA, a2.rlkA};
    bind crshA {a1.rshA, a2.rshA};
    bind crlkB {a1.rlkB, a2.rshB};
    bind crshB {a1.rshB, a2.rlkB}; //illegal
  }
}
```

Example 128—Hierarchical resource binding

Copyright © 2026 Accellera. All rights reserved.

12. Pools

Pools are used to determine possible assignment of objects to actions, and thus shape the space of legal test scenarios. *Pools* represent collections of resources, state variables, and connectivity for data flow purposes. Flow object exchange is always mediated by a pool. One action outputs an object to a pool and another action inputs it from that same pool. Similarly, actions **lock** or **share** a resource object within some pool.

Pools are structural entities instantiated under components. They are used to determine the accessibility that **action**s (see 9.2) have to flow and resource objects. This is done by binding object reference fields of action types to pools of the respective object types. Bind directives in the component scope associate resource references with a specific resource pool, state references with a specific state pool (or state variable), and buffer/stream object references with a specific data flow object pool (see 12.3).

12.1 Syntax

component_pool_declaration ::= **pool** [**[** expression **]**] type_identifier identifier **;**

Syntax 55—Instantiation of pool

In Syntax 55, *type_identifier* refers to a flow/resource object type, i.e., a **buffer**, **stream**, **state**, or **resource** struct type.

The *expression* applies only to pools of resource type; it specifies the number of resource instances in the pool. If omitted, the size of the resource pool defaults to `1`.

The following also apply:

a) The execution semantics of a pool are determined by its object type.

b) A pool of **state** type can hold one object at any given time, a pool of **resource** type can hold up to the given maximum number of unique resource objects throughout a scenario, and a pool of **buffer** or **stream** type is not restricted in the number of objects at a given time or throughout the scenario.

12.2 Example

Example 129 demonstrates how to declare a pool.

```
buffer data_buff_s {
  rand mem_segment_s seg;
};
resource channel_s {...};
component dmac_c {
  pool data_buff_s buff_p;
  ...
  pool [4] channel_s chan_p;
}
```

Example 129—Declaration of pool

Copyright © 2026 Accellera. All rights reserved.

12.3 Static pool binding directive

Every action executes in the context of a single component instance, and every object resides in some pool. Multiple actions may execute concurrently, or over time, in the context of the same component instance, and multiple objects may reside concurrently, or over time, in the same pool. Actions of a specific component instance output objects to or input objects from a specific pool. Actions of a specific component instance can only be assigned a resource of a certain pool.

Static **bind** directives determine which pools are accessible to the actions' object references under which component instances (see Syntax 56). Binding is done relative to the component sub-tree of the component type in which the **bind** directive is applied. See also 17.1.

12.3.1 Syntax

```
object_bind_stmt ::= bind hierarchical_id object_bind_item_or_list ;
object_bind_item_or_list ::=
    object_bind_item_path
  | { object_bind_item_path { , object_bind_item_path } }
object_bind_item_path ::= { component_path_elem . } object_bind_item
component_path_elem ::= component_identifier { [ domain_open_range_list ] }
object_bind_item ::=
    action_type_identifier . identifier { [ domain_open_range_list ] }
  | *
```

Syntax 56—Static bind directives

Pool binding can take one of two forms:

- *Explicit binding:* associating a pool with a specific object reference field (input/output/resource-claim) of an action type under a component instance or one or more elements of a component instance array.
- *Default binding:* associating a pool generally with a component instance sub-tree, or array of component instances, by object type.

The following also apply:

a) Components (and arrays thereof) and pools are identified with a relative instance path expression. A specific object reference field is identified with the component instance path expression, followed by an action-type name and field name, separated by dots (.).

b) Default binding can be specified for an entire sub-tree by using a wildcard instead of specific paths. When referring to an entire array, the array may be referred to by name, without needing to specify the range of elements in brackets ("[]"). Also, in case of a multidimensional array, it is possible to specify a range only to some dimensions.

c) Explicit binding always takes precedence over default bindings.

d) Conflicting explicit bindings for the same object reference field shall be illegal.

e) If multiple bindings apply to the same object reference field, the **bind** directive in the context of the top-most component instance takes precedence (i.e., the order of default binding resolution is top-down).

f) Applying multiple default bindings to the same object reference field(s) from the same component shall be illegal.

Copyright © 2026 Accellera. All rights reserved.

g) When binding object reference fields to a pool, the object and the pool shall be of the exact same type. Thus, it shall be illegal to bind an object of a derived type to a pool of its base type, or vice versa.

12.3.2 Examples

Example 130 illustrates default binding pools.

In these examples, the `buff_p` pool of `data_buff_s` objects is bound using the wildcard specifier (`{*}`). Because the **bind** statement is applied in the context of component `dma_c`, the `buff_p` pool is bound to all component instances and actions defined in `dma_c` (i.e., component instances `dmas1` and `dmas2`, and action `mem2mem_a`). Thus, the `in_data` input and `out_data` output of the `mem2mem_a` action share the same `buff_p` pool. The `chan_p` pool of `channel_s` resources is bound to the two instances.

```
struct mem_segment_s {...};
buffer data_buff_s {
  rand mem_segment_s seg;
};
resource channel_s {...};
component dma_sub_c {
  ...
};
component dma_c {
  dma_sub_c dmas1, dmas2;
  pool data_buff_s buff_p;
  bind buff_p {*};
  pool [4] channel_s chan_p;
  bind chan_p {dmas1.*, dmas2.*};
  action mem2mem_a {
    input data_buff_s in_data;
    output data_buff_s out_data;
    ...
  };
};
```

Example 130—Static binding

Example 131 illustrates the binding of pools to arrays of components. Each declared pool is of a different type, each of which will be bound to a different subset of the array of `mem_c` components.

Copyright © 2026 Accellera. All rights reserved.

```
component mem_c {...}

component top_c {
    mem_c mem[4];

    pool mbuf  mbuf_p;
    pool mbuf2 mbufA_p;
    pool mbuf3 mbufB_p;
    pool mbuf4 mbufC_p;

    bind mbuf_p  mem.*;       // All elements of the array
    bind mbufA_p mem[0..2].*; // Explicit range
    bind mbufB_p mem[1..].*;  // Up to the top element of the array
    bind mbufC_p mem[2,3].*;  // Explicit array element(s)
    ...
}
```

Example 131—Binding of pools to array of components

Example 132 illustrates the two forms of binding:, explicit and default. Action `power_transition_a`'s input and output are both associated with the context component's (`graphics_c`) state object pool. However, action `observe_same_power_state_a` has two inputs, each of which is explicitly associated with a different state object pool, the respective sub-component state variable. The `channel_s` resource pool is instantiated under the multimedia subsystem and is shared between the two engines.

Copyright © 2026 Accellera. All rights reserved.

```
state power_state_s { rand int in [0..4] level; }
resource channel_s {}
component graphics_c {
   pool power_state_s power_state_var;
   bind power_state_var *; // accessible to all actions under this
                           // component (specifically power_transition's
                           // input/output)
   action power_transition_a {
      input power_state_s curr;  //current state
      output power_state_s next; //next state
      lock channel_s chan;
   }
}
component my_multimedia_ss_c {
   graphics_c gfx0;
   graphics_c gfx1;
   pool [4] channel_s channels;
   bind channels {gfx0.*,gfx1.*};// accessible by default to all actions
                                 // under these component sub-trees
                                 // (specifically power_transition's chan)
   action observe_same_power_state_a {
      input power_state_s gfx0_state;
      input power_state_s gfx1_state;
      constraint gfx0_state.level == gfx1_state.level;
   }
   // explicit binding of the two power state variables to the
   // respective inputs of action observe_same_power_state_a
   bind gfx0.power_state_var observe_same_power_state_a.gfx0_state;
   bind gfx1.power_state_var observe_same_power_state_a.gfx1_state;
}
```

Example 132—Pool binding

Copyright © 2026 Accellera. All rights reserved.

In Example 133, there is a observe_same_power_state_a **action** type with an array of 2 input **state** objects. Action power_transition_a will cause at least one inferred instance to bind with the respective observe_same_power_state_a action's object for each one of the graphics_c component instances. Using explicit pool bind statements, each element in the object array of observe_same_power_state_a is bound to a different pool.

```
state power_state_s {
   rand int in [0..4] level;
   constraint initial -> level == 0;
}

// graphics component with power state
component graphics_c {
   pool power_state_s power_state_var;
   bind power_state_var *; // accessible to all actions under this
                           // component (specifically power_transition's
                           // input/output)
   action power_transition_a {
      input  power_state_s curr; //current state
      output power_state_s next; //next state
   }
}

component my_multimedia_ss_c {
   graphics_c gfx[2];

   action observe_same_power_state_a {
      rand int in [1..4] observed_level;

      input power_state_s gfx_state[2];
      constraint { foreach (s: gfx_state) {
        s.level == observed_level;
      }}
   }
   // explicit binding of the two power state variables to the
   // respective inputs of action observe_same_power_state_a
   bind gfx[0].power_state_var observe_same_power_state_a.gfx_state[0];
   bind gfx[1].power_state_var observe_same_power_state_a.gfx_state[1];
}
```

Example 133—Multiple state pools of the same state type

Copyright © 2026 Accellera. All rights reserved.

12.4 Resource pools and the instance_id attribute

Each object in a resource pool has a unique **instance_id** value, ranging from 0 to the pool's size − 1. Two actions that reference a resource object with the same **instance_id** value in the same pool are referencing the same resource object. See also 13.1.

For example, in Example 134, action transfer is locking two kinds of resources: channel_s and cpu_core_s. Because channel_s is defined under component dma_c, each dma_c instance has its own pool of two channel objects. Within action par_dma_xfers, the two transfer actions can be assigned the same channel **instance_id** because they are associated with different dma_c instances. However, these same two actions shall be assigned a different cpu_core_s object, with a different **instance_id**, because both dma_c instances are bound to the same resource pool of cpu_core_s objects defined under **pss_top** and they are scheduled in parallel. The **bind** directive designates the pool of cpu_core_s resources is to be utilized by both instances of the dma_c component.

```
resource cpu_core_s {}
component dma_c {
  resource channel_s {}
  pool[2] channel_s channels;
  bind channels {*}; // accessible to all actions
                     // under this component (and its sub-tree)
  action transfer {
    lock channel_s chan;
    lock cpu_core_s core;
    }
}
component pss_top {
  dma_c dma0,dma1;
  pool[4] cpu_core_s cpu;
  bind cpu {dma0.*, dma1.*};// accessible to all actions
                            // under the two sub-components
  action par_dma_xfers {
    dma_c::transfer xfer_a;
    dma_c::transfer xfer_b;

    constraint xfer_a.comp != xfer_b.comp;
    constraint xfer_a.chan.instance_id==xfer_b.chan.instance_id; //OK
   constraint xfer_a.core.instance_id==xfer_b.core.instance_id; //conflict!
   activity {
      parallel {
         xfer_a;
         xfer_b;
      }
   }
  }
}
```

Example 134—Resource object assignment

12.5 Pool of states and the initial attribute

Each pool of a **state** type contains exactly one state object at any given point in time throughout the execution of the scenario. A state pool serves as a state variable instantiated in the context component. Actions outputting to a state pool can be viewed as transitions in a finite state machine. See also 13.1.

Copyright © 2026 Accellera. All rights reserved.

Prior to execution of an action that outputs a state object to the pool, the pool contains the initial object. The **`initial`** flag is *true* for the initial object and *false* for all other objects subsequently residing in the pool. The initial state object is overwritten by the first state object (if any) which is output to the pool. The initial object is only input by actions that are scheduled before any action that outputs a state object to the same pool.

Consider, for example, the code in Example 135. The action `codec_c::configure` has an `UNKNOWN` mode as its configuration state precondition, due to the constraint on its input `prev_conf`. Because it outputs a new state object with a different mode value, there can only be one such action per `codec` component instance (unless another action, not shown here, sets the mode back to `UNKNOWN`).

```
enum codec_config_mode_e {UNKNOWN, A, B}
component codec_c {
  state configuration_s {
    rand codec_config_mode_e mode;
    constraint initial -> mode == UNKNOWN;
  }
  pool configuration_s config_var;
  bind config_var *;
  action configure {
    input configuration_s prev_conf;
    output configuration_s next_conf;
    constraint prev_conf.mode == UNKNOWN && next_conf.mode in [A, B];
  }
}
```

Example 135—State object binding

Copyright © 2026 Accellera. All rights reserved.

13. Randomization

Scenario properties can be expressed in PSS declaratively, as algebraic constraints over attributes of scenario entities.

a) There are several categories of **struct** and **action** fields.
   1) *Random attribute field* - a field of a plain-data type (e.g., **bit**) that is qualified with the **rand** keyword.
   2) *Non-random attribute field* - a field of a plain-data type (e.g., **int**) that is not qualified with the **rand** keyword.
   3) *Sub-action field* - a field of an action type or a plain-data type that is qualified with the **action** keyword.
   4) *Input/output flow object reference field* - a field of a flow object type that is qualified with the **input** or **output** keyword.
   5) *Resource claim reference field* - a field of a resource object type that is qualified with the **lock** or **share** keyword.

b) Constraints may shape every aspect of the scenario space. In particular:
   1) Constraints are used to determine the legal value space within the type domain for attribute fields of actions.
   2) Constraints affect the legal assignment of resources to actions and, consequently, the scheduling of actions.
   3) Constraints may restrict the possible binding of action inputs to action outputs, and, thus, possible action inferences from partially specified scenarios.
   4) Constraints determine the association of actions with context component instances.
   5) Constraints may be used to specify all of the above properties in a specific context of a higher level activity encapsulated via a compound action.
   6) Constraints may also be applied also to the operands of control flow statements—determining loop count and conditional branch selection.

Constraints are typically satisfied by more than just one specific assignment. There is often room for randomness or the application of other considerations in selecting values. The process of selecting values for scenario variables is called *constrained randomization* or simply *randomization*.

Randomized values of variables become available in the order in which they are used in the execution of a scenario, as specified in activities. This provides a natural way to express and reason about the randomization process. It also guarantees values sampled from the environment and fed back into the PSS domain during the generation and/or execution have clear implications on subsequent evaluation. However, this notion of ordering in variable randomization does not introduce ordering into the constraint system—the solver is required to look ahead and accommodate for subsequent constraints.

13.1 Algebraic constraints

13.1.1 Member constraints

PSS supports two types of constraint blocks (see Syntax 57) as **action**/**struct** members: fixed constraints that always hold and generic constraints that only hold when they are referenced by the user by traversing them in an activity (see 13.4.11) or referencing them inside a constraint. A generic constraint associates a name and, optionally parameters, with a constraint. Generic constraints only apply once they are referenced, directly or indirectly, from a fixed constraint.

NOTE—*dynamic* constraints are deprecated. Their functionality is replaced by a generic constraint with no parameters.

Copyright © 2026 Accellera. All rights reserved.

13.1.1.1 Syntax

```
constraint_declaration ::=
    constraint constraint_set
  | [ dynamic ] constraint identifier constraint_block
  | generic_constraint_declaration
constraint_set ::=
    constraint_body_item
  | constraint_block
constraint_block ::= { { constraint_body_item } }
constraint_body_item ::=
    expression_constraint_item
  | foreach_constraint_item
  | forall_constraint_item
  | if_constraint_item
  | implication_constraint_item
  | unique_constraint_item
  | default hierarchical_id == constant_expression ;
  | default disable hierarchical_id ;
  | dist_directive
  | constraint_body_compile_if
  | soft_constraint_item
  | annotation_usage constraint_body_item
  | stmt_terminator
```

Syntax 57—Member constraint declaration

13.1.1.2 Examples

Example 136 declares a fixed constraint block, while Example 137 declares a generic constraint block. In the case of the fixed constraint, the name is optional.

```
action A {
  rand bit[32]    addr;

  constraint addr_c {
    addr == 0x1000;
  }
}
```

Example 136—Declaring a static constraint

Copyright © 2026 Accellera. All rights reserved.

```
action B {
  action bit[32]    addr;

  constraint dyn_addr1_c() {
    addr in [0x1000..0x1FFF];
  }

  constraint dyn_addr2_c() {
    addr in [0x2000..0x2FFF];
  }
}
```

Example 137—Declaring a generic constraint

13.1.2 Generic constraints

13.1.2.1 Syntax

```
generic_constraint_declaration ::=
  generic_constraint_bool
  | generic_constraint_value
generic_constraint_bool ::=
  [static] constraint identifier generic_constraint_params constraint_set
generic_constraint_value ::=
  [static] constraint generic_constraint_data_type identifier
    generic_constraint_params expression_constraint_item ;
generic_constraint_params ::=
  ([generic_constraint_param {, generic_constraint_param}])
generic_constraint_param ::=
  [const] generic_constraint_data_type identifier
generic_constraint_data_type ::=
    numeric
    | data_type
```

Syntax 58—Generic constraint declaration

Generic constraints support definition of reusable constraints that are defined in terms of a set of input expressions. Once defined, the expression or statements can be reused in multiple contexts with different expressions provided for the parameters.

Generic constraints can be declared in struct, action, and component scopes, as well as in package scopes where they are always static.

The following also apply:

a) Generic constraints may not declare local variables

b) Generic constraints with a result type are defined as a single expression, and may be used in any position within a constraint where an expression of that type is legal.

Copyright © 2026 Accellera. All rights reserved.

c) When a generic constraint in a struct, action, or component scope shadows a constraint declared in a base type, the return and parameter types shall match.

d) Recursive references are supported, provided that recursion is gated by an expression that does not involve randomization.

13.1.2.2 Examples

```
action send_pkt {
  rand bit[16] pkt_sz;

  constraint pkt_sz_c {pkt_sz > 0;}

  constraint interesting_sz_c {small_pkt_c() || jumbo_pkt_c();}

  constraint small_pkt_c() {pkt_sz <= 100;}
  constraint jumbo_pkt_c() {pkt_sz > 1500;}
}

action scenario {
  activity {
    // Send a packet with size in [1..100, 1501..65535]
    do send_pkt;
    // Send a small packet with a directly-specified in-line constraint
    do send_pkt with {pkt_sz <= 100;};
    // Send a small packet by referencing a generic constraint
    do send_pkg with {small_pkt_c();};
  }
}
```

Example 138—Referencing a generic constraint inside a static constraint

Example 138 shows a generic constraint inside a fixed constraint. In the examples, the `send_pkt` action sends a packet of a random size. The fixed constraint `pkt_sz_c` ensures the packet is of a legal size and the two generic constraints, `small_pkt_c` and `jumbo_pkt_c`, specialize the packet size to be small or large, respectively. The fixed constraint `interesting_sz_c` restricts the size to be either `<=100` for `small_pkt_c` or `>1500` for `jumbo_pkt_c`.

```
constraint numeric max(numeric a, numeric b)
  (a < b)?b:a;

struct S {
  rand bit[8] j, k, l;
  // Equivalent to:
  // j == (k < l) ? l : k;
  constraint j == max(k,l);
}
```

Example 139—Value-yielding generic constraint

The generic constraint in Example 139 yields the maximum of its two parameters. Because the constraint operates on numeric, it can be used for both floating-point and integer expressions.

Copyright © 2026 Accellera. All rights reserved.

```
constraint gt_a_list_elem(
  bit[64] val,
  list<bit[64]> l) {
  gt_a_list_elem_inner(val, l, 0);
}
constraint gt_a_list_elem_inner(
  bit[64] val,
  list<bit[64]> l,
  const bit[64] idx) {
  if (idx+1 < l.size()) {
    (val > l[idx]) || gt_a_list_elem_inner(val, l, idx+1);
  } else {
    (val > l[idx]);
  }
}
```

Example 140—Generic recursive constraint

Example 140 shows a generic constraint that ensures that the specified value expression is greater than at least one of the values in the array. In other words, it forms a constraint of the form:

```
((val > list[0]) || (…) || (val > list[SZ-1]))
```

The inner constraint *gt_a_list_elem_inner* tracks expansion of the OR terms by using a state variable. Recursion is controlled by an invariant condition—a condition that involves no random expressions.

13.1.3 Constraint inheritance

As discussed in 17.1, an **action**/**struct** subtype has all of the constraints that are declared in the context of its supertype or that are inherited by the supertype. Unnamed fixed constraints in a subtype are added to all other constraints. A named fixed or generic **constraint** in a subtype *shadows* (masks) a constraint of the same name from the supertype. Constraint inheritance applies in the same way to fixed constraints and generic constraints.

Example 141 illustrates a simple case of constraint inheritance and shadowing. Instances of **struct** `corrupt_data_buff` satisfy the unnamed constraint of `data_buff` based on which `size` is in the range `1` to `1024`. Additionally, `size` is greater than `256`, as specified in the subtype. Finally, per constraint `size_align` as specified in the subtype, `size` divided by `4` has a reminder of `1`.

```
buffer data_buff {
    rand int size;
    constraint size in [1..1024];
    constraint size_align { size%4 == 0; }       // 4-byte aligned
}

buffer corrupt_data_buff : data_buff {
    constraint size_align { size%4 == 1; }       // alignment 1 byte off
    constraint corrupt_data_size { size > 256; } // additional constraint
}
```

Example 141—Inheriting and shadowing constraints

Copyright © 2026 Accellera. All rights reserved.

13.1.4 Action traversal in-line constraints

Constraints on sub-action data attributes can be in-lined directly in the context of an *action traversal statement* in the **activity** clause (for syntax and other details, see 11.3.1).

In the context of in-line constraints, attribute field paths of the traversed sub-action can be accessed without the sub-action field qualification. Fields of the traversed sub-action take precedence over fields of the containing action. Other attribute field paths are evaluated in the context of the containing action. In cases where the containing-action fields are shadowed (masked) by fields of the traversed sub-action, they can be explicitly accessed using the built-in variable **`this`**. In particular, fields of the context component of the containing action shall be accessed using the prefix path **`this.comp`** (see also Example 143).

If a sub-action field is traversed uniquely by a single traversal statement in the **activity** clause, in-lining a constraint has the same effect as declaring the same member constraint on the sub-action field of the containing action. In cases where the same sub-action field is traversed multiple times, in-line constraints apply only to the specific traversal in which they occur.

Unlike member constraints, in-line constraints are evaluated in the specific scheduling context of the *action traversal statement*. If attribute fields of sub-actions other than the one being traversed occur in the constraint, these sub-action fields shall have already been traversed in the activity. In cases where a sub-action field has been traversed multiple times, the most recently selected values are considered.

Example 142 illustrates the use of in-line constraints. The traversal of `a3` is illegal, because the path `a4.f` occurs in the in-line constraint, but `a4` has not yet been traversed at that point. Constraint `c2`, in contrast, equates `a1.f` with `a4.f` without having a specific scheduling context, and is, therefore, legal and enforced.

```
action A {
  rand bit[4]   f;
};

action B {
  A a1, a2, a3, a4;

  constraint c1 { a1.f in [8..15]; };
  constraint c2 { a1.f == a4.f; };

  activity {
    a1;
    a2 with {
      f in [8..15]; // same effect as constraint c1 has on a1
    };
    a3 with {
      f == a4.f;    // illegal: a4.f unresolved at this point
    };
    a4;
  }
};
```

Example 142—Action traversal in-line constraint

Copyright © 2026 Accellera. All rights reserved.

Example 143 illustrates different name resolutions within an in-line **with** clause.

```
component subc {
    action A {
        rand int f;
        rand int g;
    }
}

component top {
    subc sub1, sub2;
    action B {
        rand int f;
        rand int h;
        subc::A a;

        activity {
            a with {
                f < h;        // sub-action's f and containing action's h
                g == this.f; // sub-action's g and containing action's f
                comp == this.comp.sub1;
                              // sub-action's component is sub-component
                              // 'sub1' of the parent action's component
            };
        }
    }
}
```

Example 143—Name resolution inside **with** *constraint block*

13.1.5 Logical expression constraints

A logical (Boolean) constraint can be used to specify a constraint. Syntax 59 shows the syntax for an expression constraint.

13.1.5.1 Syntax

```
expression_constraint_item ::= expression ;
```

Syntax 59—Expression constraint

expression may be any logical expression. The constraint is satisfied if the expression evaluates to *true*.

13.1.6 Implication constraints

Conditional constraints can be specified using the *implication* operator (**->**). Syntax 60 shows the syntax for an implication constraint.

13.1.6.1 Syntax

```
implication_constraint_item ::= expression -> constraint_set
```

Syntax 60—Implication constraint

Copyright © 2026 Accellera. All rights reserved.

expression may be any logical expression. *constraint_set* represents any valid constraint or an unnamed constraint set.

The following also apply:

a) The Boolean equivalent of the implication operator `a -> b` is `(!a || b)`. This states that if the *expression* is *true*, all of the constraints in *constraint_set* shall be satisfied. In other words, if the *expression* is *true*, then the random values generated are constrained by the constraint set. Otherwise, the random values generated are unconstrained.

b) The implication constraint is bidirectional.

13.1.6.2 Example

Consider Example 144. Here, `b` is forced to have the value `1` whenever the value of the variable `a` is greater than `5`. However, since the constraint is bidirectional, if `b` has the value `1`, then the evaluation expression `(!(a>5) || (b==1))` is *true*, so the value of `a` is unconstrained. Similarly, if `b` has a value other than `1`, `a` is `<= 5`.

```
struct impl_s {
  rand bit[8]     a, b;

  constraint ab_c {
    (a > 5) -> b == 1;
  }
}
```

Example 144—Implication constraint

13.1.7 If-else constraints

Conditional constraints can be specified using the **if** and **if-else** constraint statements.

Syntax 61 shows the syntax for an **if-else** constraint.

13.1.7.1 Syntax

```
if_constraint_item ::= if ( expression ) constraint_set [ else constraint_set ]
```

Syntax 61—Conditional constraint

expression may be any logical expression. *constraint_set* represents any valid constraint or an unnamed constraint set.

The following also apply:

a) If the *expression* is *true*, all of the constraints in the first *constraint_set* shall be satisfied; otherwise, all the constraints in the optional **else** *constraint_set* shall be satisfied.

b) Constraint sets may be used to group multiple constraints.

c) Just like *implication* (see 13.1.6), *if-else* style constraints are bidirectional.

13.1.7.2 Examples

In Example 145, the value of `a` constrains the value of `b` and the value of `b` constrains the value of `a`.

Copyright © 2026 Accellera. All rights reserved.

Attribute `a` cannot take the value `0` because both alternatives of the **if-else** constraint preclude it. The maximum value for attribute `b` is `4`, since in the `if` alternative it is `1` and in the `else` alternative it is less than `a`, which itself is `<= 5`.

In evaluating the constraint, the `if`-clause evaluates to `!(a>5) || (b==1)`. If `a` is in the range `{1,2,3,4,5}`, then the `!(a>5)` expression is *true*, so the `(b==1)` constraint is ignored. The `else`-clause evaluates to `!(a<=5)`, which is *false*, so the constraint expression `(b<a)` is *true*. Thus, `b` is in the range `{0..(a-1)}`. If `a` is `2`, then `b` is in the range `{0,1}`. If `a > 5`, then `b` is `1`.

However, if `b` is `1`, the `(b==1)` expression is *true*, so the `!(a>5)` expression is ignored. At this point, either `!(a<=5)` or `a > 1`, which means that `a` is in the range `{2,3, … 255}`.

```
struct if_else_s {
  rand bit[8]     a, b;

  constraint ab_c {
    if (a > 5) {
      b == 1;
    } else {
      b < a;
    }
  }
}
```

Example 145—If constraint

Copyright © 2026 Accellera. All rights reserved.

13.1.8 Foreach constraints

Elements of collections can be iteratively constrained using the **foreach** constraint.

Syntax 62 shows the syntax for a **foreach** constraint.

13.1.8.1 Syntax

```
foreach_constraint_item ::=
  foreach ( [ iterator_identifier : ] expression [ [ index_identifier ] ] ) constraint_set
```

Syntax 62—Foreach constraint

constraint_set represents any valid constraint or an unnamed constraint set.

The following also apply:

a) *expression* shall be of a collection type (i.e., **array**, **list**, **map** or **set**), including fixed-sized arrays of *monitor elements*.

b) All of the constraints in *constraint_set* shall be satisfied for each of the elements in the collection specified by *expression*.

c) *iterator_identifier* specifies the name of an iterator variable of the collection element type. Within *constraint_set*, the iterator variable, when specified, is an alias to the collection element of the current iteration.

d) *index_identifier* specifies the name of an index variable. Within *constraint_set*, the index variable, when specified, corresponds to the element index of the current iteration.

   1) For **array**s and **list**s, the index variable shall be a variable of type **int**, ranging from `0` to one less than the size of the collection variable.

   2) For **map**s, the index variable shall be a variable of the same type as the **map** keys, and range over the values of the keys.

   3) For **set**s, an index variable shall not be specified.

e) Both the index and iterator variables, if specified, are implicitly declared within the **foreach** scope and limited to that scope. Regular name resolution rules apply when the implicitly declared variables are used within the **foreach** body. For example, if there is a variable in an outer scope with the same name as the index variable, that variable is shadowed (masked) by the index variable within the **foreach** body. The index and iterator variables are not visible outside the **foreach** scope.

f) Either an index variable or an iterator variable or both shall be specified. For a **set**, an iterator variable shall be specified, but not an index variable.

13.1.8.2 Example

Example 146 shows an iterative constraint that ensures that the values of the elements of a fixed-size array increment.

Copyright © 2026 Accellera. All rights reserved.

```
struct foreach_s {
  rand bit[10]   fixed_arr[10];

  constraint fill_arr_elem_c {
    foreach (fixed_arr[i]) {
      if (i > 0) {
        fixed_arr[i] > fixed_arr[i-1];
      }
    }
  }
}
```

Example 146—Foreach iterative constraint

13.1.9 Forall constraints

The **forall** constraint is used to apply constraints to all instances of a specific type within the instance subtree in which the constraint is placed.

Syntax 63 shows the syntax for a **forall** constraint.

13.1.9.1 Syntax

```
forall_constraint_item                                                    ::=
  forall ( iterator_identifier : type_identifier [ in ref_path ] ) constraint_set
```

Syntax 63—Forall constraint

type_identifier specifies the type of the entity (**action**, **struct**, **stream**, **buffer**, **state**, **resource**) to which the constraint applies. *iterator_identifier* can be used inside *constraint_set* as an alias to each instance, much like the *iterator_identifier* in a **foreach** constraint is an alias to each element in the collection (see 13.1.8). *ref_path* is optionally used to restrict the constraint's scope of application to a certain instance subtree.

The following also apply:

a) All of the constraints in *constraint_set* shall be satisfied for every instance of the specified type in the **forall** constraint's application scope.

b) When *ref_path* is omitted, the application scope is the subtree of the constraint's enclosing scope:

   1) In the case of a member (type-level) fixed constraint, its application scope includes all of the context type's fields (attributes, object references), and in the case of a compound action, also its entire activity.

   2) In the case of an in-line **with** constraint (see 13.1.4), its application scope is the traversed sub-action's fields and, if compound, also its entire activity.

   3) In the case of an activity constraint statement or the activation of a named generic constraint, the application scope is the activity scope immediately enclosing the activity statement.

c) When *ref_path* is specified, the application scope is the subtree under the entity (**action**, object, or **struct**) designated by *ref_path*.

d) The **forall** constraint applies to sub-actions within its application scope regardless of whether they are traversed using an action handle or anonymously.

Copyright © 2026 Accellera. All rights reserved.

13.1.9.2 Examples

Example 147 demonstrates the use of a **forall** constraint in a compound action, constraining sub-actions traversed directly and indirectly under its activity (case b.1 above). Action `entry` places a constraint on all instances of action `A`, relating attribute `x` to its own attribute `ax_limit`. The constraint does not apply to an attribute of sub-action `B` by the same name.

```
action A {
  rand int in [0..9] x;
};

action B {
  rand int in [0..9] x;
};

action C {
  A a;
  B b;
  activity {
    schedule {
      a; b;
    }
  }
};

action entry {
  rand int in [0..9] ax_limit;
  A a;
  C c;
  constraint {
    forall (a_it: A) {
      a_it.x <= ax_limit;
    }
  }
  activity {
    a; c;
  }
};
```

Example 147—Forall constraint

The **forall** constraint in Example 147 is equivalent to the corresponding constraint on each path to an action handle of type `A`. Hence, action `entry` in Example 147 can be rewritten in the way shown in Example 148.

Copyright © 2026 Accellera. All rights reserved.

```
action entry {
  rand int in [0..9] ax_limit;
  A a;
  C c;
  constraint {
    a.x <= ax_limit;
    c.a.x <= ax_limit;
  }
  activity {
    a; c;
  }
};
```

Example 148—Rewrite of forall constraint in terms of explicit paths

Example 149 demonstrates the use of **forall** constraints in two different contexts inside an activity. The first is an in-line **with** constraint item (case b.2 above), applying to all instances of type `A` under action `C` that is being traversed in this statement. The second is an activity constraint statement (case b.3 above). It applies to all instances of type `A` in the immediately enclosing activity scope – in this case the **parallel** statement. Hence this constraint applies to action `A` in the first **parallel** branch, and to all actions of type `A` under action `C` in the second **parallel** branch.

```
action entry {
  activity {
    do C with {
      forall (a_it: A) {
        a_it.x == 1;
      }
    }
    parallel {
      do A;
      do C;
      constraint forall (a_it: A) {
        a_it.x in [2, 4];
      }
    }
  }
};
```

Example 149—Forall constraint in different activity scopes

Example 150 demonstrates the use of a **forall** constraint item in a generic constraint under an action. The generic constraint is activated from above for one traversal of that action, and not for the other. In this case, `A`'s attributes `s1.x` and `s2.x` may be randomized to the value `0xff` in the first execution of `B`, but not in the second.

Copyright © 2026 Accellera. All rights reserved.

```
struct S {
  rand bit[8] x;
};

action A {
  rand S s1, s2;
};

action B {
  constraint c1() {
    forall (it: S) { it.x != 0xff; }
  }
  activity { do A; }
};

action entry {
  activity {
    do B;
    do B with { c1; };
  }
};
```

Example 150—Forall constraint item in a generic constraint

13.1.10 Unique constraints

The **unique** constraint causes unique values to be selected for each element in the specified set.

Syntax 64 shows the syntax for a **unique** constraint.

13.1.10.1 Syntax

```
unique_constraint_item ::= unique unique_constraint_argument ;
unique_constraint_argument ::=
        hierarchical_id  [ [ array_slice ] ]
        | { hierarchical_id_list }
hierarchical_id_list ::= hierarchical_id { , hierarchical_id }
```

Syntax 64—Unique constraint

The set of elements can be specified in one of the two forms.

In the first form, a single array or list attribute is specified, possibly with a slice of indexes that designates a sub-array or a sub-list. If no slice is specified, the values of the array or list elements shall be unique. If a slice is specified, the values of the elements starting from the left slice index (if it is non-negative) and/or ending at the right slice index (if it is lower than the array/list size) shall be unique. The slice indexes themselves can be random. If the left slice index is greater than the right slice index, or the left slice index is greater than or equal to the array/list size, or the right slice index is negative, the slice is considered empty, and the **unique** constraint is satisfied vacuously regardless of the array/list element values.

In the second form, an explicit set of attributes, enclosed in curly braces, is specified. The values of these attributes shall be unique.

Copyright © 2026 Accellera. All rights reserved.

13.1.10.2 Examples

Example 151 forces the solver to select unique values for the random attribute fields `A`, `B`, and `C`. The **unique** constraint is equivalent to the following constraint statement: `((A != B) && (A != C) && (B != C))`.

```
struct my_struct {
    rand bit[4] in [0..12] A, B, C;
    constraint unique_abc_c {
        unique {A, B, C};
    }
}
```

Example 151—Unique constraint

Example 152 forces the solver to select unique values for the elements of the array `my_arr`. All the array element values shall be different.

```
struct my_struct {
    rand array<bit[4],5> my_arr;
    constraint unique_arr{
        unique my_arr;
    }
}
```

Example 152—Unique constraint on an array

Example 153 forces the solver to select unique values for a list slice. The list element values between the indexes `M` and `N` shall be different, whereas the values of `M` and `N` shall also be randomized.

Note that the constraint `M <= N` is essential here. Without this constraint, the effect of the entire set of constraints would be different: any solution where `M > N` would satisfy the **unique** constraint vacuously, regardless of the values in `my_list`.

Copyright © 2026 Accellera. All rights reserved.

```
struct my_struct {
    rand list<int> my_list;
    int my_list_size;
    rand int M, N;
    constraint {
        M <= N;
        M < my_list_size;
        N >= 0;
    }
    constraint unique_list_slice {
        unique my_list[M..N];
    }
    exec pre_solve {
        // initialize the list size
        repeat (5) {
            my_list.push_back(0);
        }
        my_list_size = my_list.size();
    }
}
```

Example 153—Unique constraint on a list slice

13.1.11 Default value constraints

A default value constraint determines the value of an attribute, unless explicitly disabled for that specific attribute from its direct or indirect containing type. Default value constraints may only take the form of equality of the attribute to a constant expression. Disabling a default value is done with the **default disable** constraint form.

13.1.11.1 Syntax

```
constraint_body_item ::=
  ...
  | default hierarchical_id == constant_expression ;
  | default disable hierarchical_id ;
  | ...
```

Syntax 65—Default constraints

The following also apply:

a) A **default** value constraint has the same semantics as the corresponding equality constraint, unless explicitly disabled. The equality shall hold, and conflict with other constraints shall be flagged as a contradiction.

b) A **default disable** constraint is a directive to remove default constraints on the designated attribute, if any are specified.

c) *hierarchical_id* for both **default** and **default disable** constraints shall be a random attribute (a field with **rand** modifier). It shall be an error to apply a **default** constraint on a non-**rand** attribute.

d) Multiple **default** constraints and **default disable** constraints may be applied to the same attribute, with the following precedence rules:

Copyright © 2026 Accellera. All rights reserved.

1) A constraint from a higher-level containing context overrides one from a lower-level containing context.
2) A constraint from a derived type context overrides one from a base type context.
3) A constraint overrides another in the same type context if it occurs later in the code.

e) **default** value constraints and **default disable** constraints may be applied to an attribute of an aggregate data type. The semantics in this case are equivalent to applying the corresponding constraints to all the **rand** scalar attributes it comprises. In particular, applying a **default disable** constraint to an attribute of an aggregate data type disables **default** value constraints on all attributes under it.

f) **default** and **default disable** constraints may not be conditioned on non-constant expressions.

g) **default** and **default disable** constraints may not be used under generic constraints (constraints prefixed with the **generic** modifier).

h) When **default** and **default disable** constraints apply to an attribute of a flow object, they are treated according to the above rules for each **input** and **output** reference separately. In other words, two **default** or **default disable** constraints that apply to the same flow object instance attribute via different **input**/**output** references, shall not override or remove each other. If the resulting **default** constraints applying to two **input** and/or **output** object references contradict each other, these two references cannot refer to the same object instance. In particular, if there is an explicit **bind** statement or equality constraint between such two object references, a contradiction shall be flagged.

i) When **default** and **default disable** constraints apply to an attribute of a resource object, they are treated according to the above rules for each **lock** or **share** reference separately. In other words, two **default** or **default disable** constraints that apply to the same resource object instance attribute via different **lock**/**share** references shall not override or remove each other. If the resulting **default** constraints applying to two **lock** or **share** object references contradict each other, these two references cannot refer to the same object instance. In particular, if there is an explicit equality constraint between such two references, a contradiction shall be flagged.

13.1.11.2 Examples

In Example 154, `my_struct` has two attributes, and a **default** value constraint on one of them. This **struct** is instantiated three times under `my_action`.

```
struct my_struct {
  rand int in [0..3] attr1;
  constraint default attr1 == 0; // (1)

  rand int in [0..3] attr2;
  constraint attr1 < attr2; // (2)
};

action my_action {
  rand my_struct s1;

  rand my_struct s2;
  constraint default s2.attr1 == 2; // (3)

  rand my_struct s3;
  constraint default disable s3.attr1; // (4)
  constraint s3.attr1 > 0; // (5)
};
```

Example 154—Use of default value constraints

Copyright © 2026 Accellera. All rights reserved.

When randomizing `my_action`, `s1.attr1` is resolved to `0` because of constraint (1), and `s1.attr2` is randomized in the domain `1..3` because of constraint (2). `s2.attr1` is resolved to `2`, because constraint (3) overrides constraint (1), and `s2.attr2` is resolved to 3 because of constraint (2). Within `s3`, constraint (1) was disabled by (4), and has no effect. Due to constraints (2) and (5), `s3.attr1` is randomized in the domain `1..2` and `s3.attr2` in the domain `2..3` such that `s3.attr1` is less than `s3.attr2`.

In Example 155 below, two attributes of `my_action` have **default** value constraints. If `my_derived_action` is randomized, `attr1` is resolved to `0`, because **default** constraint (1) is disabled (3) and a different constraint is in effect (4). However, there is no consistent assignment to `attr2`, because both **default** constraint (2) and the regular constraint (5) are in effect and conflicting.

```
action my_action {
  rand int attr1;
  constraint default attr1 == -1; // (1)

  rand int attr2;
  constraint default attr2 == -1; // (2)
};

action my_derived_action : my_action {
  constraint {
    default disable attr1;        // (3)
    attr1 == 0;                   // (4) OK
  }

  constraint attr2 == 0;          // (5) contradiction!
};
```

Example 155—Contradiction with default value constraints

Example 156 below shows how **default** value constraints and **default disable** constraints apply to aggregate data types. A **default** value constraint is placed on an array as a whole (1). Under `my_action`, for instance `s1` of the struct, the default is replaced by another for a specific element (3), while the other elements retain their original default. Constraint (4) disables the default for all array elements under `s2`, and they are randomized over their full domain. Constraint (5) disables defaults of all attributes under the struct, including the 4 `arr` elements and `attr`. A subsequent constraint determines that `s3.attr` randomizes to `50`.

Copyright © 2026 Accellera. All rights reserved.

```
struct my_struct {
  rand array<int,4> arr;
  constraint default arr == {0, 10, 20, 30}; // (1)

  rand int attr;
  constraint default attr == 40; // (2)
};

action my_action {
  rand my_struct s1, s2, s3;

  constraint default s1.arr[3] == 100; // (3)

  constraint default disable s2.arr; // (4)

  constraint default disable s3; // (5)
  constraint s3.attr == 50;
};
```

Example 156—Default value constraints on compound data types

In Example 157, **default** constraints are applied to attributes of a **buffer** type. The **default** constraints applying to the different buffer references are treated separately, as follows. The resulting **default** constraints on `my_read.b_in` imply `x==2`, `y==5`. The resulting default constraints on `my_write1.b_out` imply `x==3`. The resulting **default** constraints on `my_write2.b_out` imply `x==2`, `y==10`. The resulting **default** constraints on `my_write3.b_out` leave `x` unconstrained. Therefore, `my_read.b_in` cannot bind with `my_write1.b_out` (contradicting constraints on `x`) or `my_write2.b_out` (contradicting constraints on `y`), and it shall necessarily bind with `my_write3.b_out`, with `x==2`, `y==5`.

Copyright © 2026 Accellera. All rights reserved.

```
component my_comp {
    buffer my_buff {
        rand int x, y;
        constraint default x == 2;
    }
    pool my_buff bp;
    bind bp *;
    action my_read {
        input my_buff b_in;
        constraint default b_in.y == 5;
    }
    abstract action my_write {
        output my_buff b_out;
    }
    action my_write1: my_write {
        constraint default b_out.x == 3;
          // overrides the default constraint from the buffer declaration
    }
    action my_write2: my_write {
        constraint default b_out.y == 10;
    }
    action my_write3: my_write1 {
        constraint default disable b_out.x;
          // removes the default constraint from my_write1
    }
}
component pss_top {
    my_comp c;
    action top {
        activity {
            parallel {
                do my_comp::my_write1;
                do my_comp::my_write2;
                do my_comp::my_write3;
            }
            do my_comp::my_read;
        }
    }
}
```

Example 157—Default constraints on flow objects

In Example 158, **default** constraints are applied to the attribute of a **buffer** type, and there is an explicit **bind** between two buffer references. The **default** constraints imply that `my_read.b_in` cannot bind with `my_write1.b_out`, because their resulting **default** constraints on `x` contradict each other. However, the explicit **bind** statement requires that `rd.b_in` bind with `wr1.b_out`. Therefore, this example shall flag a contradiction.

In the absence of the explicit **bind** statement, this example would result in binding `rd.b_in` with `wr2.b_out`, whose **default** constraints are aligned.

Copyright © 2026 Accellera. All rights reserved.

```
component my_comp {
    buffer my_buff {
        rand int x;
        constraint default x == 2;
    }
    pool my_buff bp;
    bind bp *;
    action my_read {
        input my_buff b_in;
    }
    abstract action my_write {
        output my_buff b_out;
    }
    action my_write1: my_write {
        constraint default b_out.x == 3;
    }
    action my_write2: my_write {
    }
}
component pss_top {
    my_comp c;
    action top {
        my_comp::my_write1 wr1;
        my_comp::my_write2 wr2;
        my_comp::my_read rd;

        bind rd.b_in wr1.b_out;

        activity {
            parallel {
                wr1;
                wr2;
            }
            rd;
        }
    }
}
```

Example 158—Default constraints on flow objects with an explicit bind

In Example 159, **default** constraints are applied to an attribute of a **buffer** type, and an explicit hierarchical bind is used to bind a compound action **input** buffer reference to its sub-action input buffer reference. However, the **default** constraints declared on those two references contradict each other. Therefore, a contradiction shall be flagged.

Copyright © 2026 Accellera. All rights reserved.

```
component pss_top {
    buffer my_buff {
        rand int x;
    }
    pool my_buff bp;
    bind bp *;
    action my_read {
        input my_buff b_in;
        constraint default b_in.x == 2;
    }
    action my_compound_read {
        input my_buff b_in;
        my_read rd1, rd2;
        constraint default b_in.x == 3;
        bind b_in rd1.b_in;
        activity {
            rd1;
            rd2;
        }
    }
    action my_write {
        output my_buff b_out;
    }
    action top {
        activity {
            do my_compound_read;
            do my_write;
        }
    }
}
```

Example 159—Default constraints on flow objects with an explicit hierarchical bind

13.1.12 Soft constraints

```
constraint_body_item ::=
  soft_constraint_item
 | …
soft_constraint_item ::= soft expression ;
```

Syntax 66—Soft constraints

The constraints described in previous sections can be thought of as “hard” constraints. These constraints always apply. If they conflict, it results in a solve contradiction.

There are many cases where some default condition shall be specified and overridden in specific situations. Default constraints provide a simple approach to doing this when the overriding condition is static in nature. Soft constraints are used when the overriding condition is dynamic (e.g., within constraint conditions).

Soft constraints express a prioritized preference for certain solutions over others. Lower-priority soft constraints are ignored when they are contradicted by higher-priority soft constraints. Hard constraints shall always be satisfied. Any soft constraint that contradicts a hard constraint is discarded.

Copyright © 2026 Accellera. All rights reserved.

Active default constraints have the same priority as hard constraints. Any soft constraint that contradicts an active default constraint is discarded.

The relative priority of soft constraints is determined by their relative position within the PSS model description. The following apply:

- Soft constraints in an inline constraint ("with" block) have the highest priority.
- Soft constraints within the scope of the same construct (e.g., constraint block, struct, action) are assigned a priority relative to their syntactic declaration order. Constraint expressions that appear later in the construct have higher priority than constraint expressions that appear earlier.
- Soft constraints in sub-attributes (e.g., struct-type attributes, action handles) have lower priority than soft constraints declared within the containing scope.
- Soft constraints in sub-attributes (e.g., struct-type attributes, action handles) are assigned a priority relative to the declaration order of the attribute within its containing scope. Soft constraints within attributes declared later in the description have higher priority than constraints declared in attributes declared earlier in the description.
- Soft constraints on flow-object outputs have higher priority than soft constraints on flow-object inputs.
- Soft constraints declared in a derived type have a higher priority than soft constraints declared in a base type.
- Soft constraints within an iterative constraint (e.g., foreach) shall have a priority relative to the iteration order. Constraints in later iterations shall have a higher priority than soft constraints in earlier iterations.
- Soft constraints within a *forall* constraint shall have a priority based on the location of the *forall* statement. Specifically, the priority shall be the same as if the constraints within the *forall* loop were expanded in-line.
- Soft constraints declared in type extensions shall have a higher priority than soft constraints declared in the initial type declaration but a lower priority than soft constraints declared in a derived type.
- The priority of soft constraints declared in type extensions is relative to the order in which type extensions are applied. The priority of soft constraints later-applied type extensions is higher than the priority of soft constraints declared in earlier-applied type extensions.

Example 160 illustrates the preceding rules:

Copyright © 2026 Accellera. All rights reserved.

```
struct S1 {
  rand bit[32] x;
  constraint soft x > 10;   // s1_1
  constraint soft x < 100;  // s1_2
}
struct S2 : S1 {
  constraint soft x in [5..9]; // s2_1
}
action A1  {
  rand bit[32] y;
  constraint soft y > 10; // a1_1
}
action A2 : A1 {
  constraint soft y in [5..9]; // a2_1
  rand S2 f1;
  rand S1 f2;
  rand S2 f3;
  constraint f1.x < f2.x;  // a2_2
}
extend action A2 {
  constraint soft y > 100; // a2_3
}
extend action A1 { // explain order
  constraint soft y > 200; // a1_2
}
action Entry {
  activity {
    do A2 with {
        soft y in [10, 20, 30]; // w_1
        soft y < f1.x;          // w_2
    }
  }
}
```

Example 160—Soft-constraint priority

The relative priority of the soft constraints (highest to lowest) in Example 160 is:

- `w_2`: In-line constraints have higher priority than type-level constraints
- `w_1`
- `a2_3`: Constraints in type extension have higher priority than initial definition constraints
- `a2_2`: Constraints declared later in a type have priority over those declared earlier
- `a2_1`
- `f3.s2_1`: Constraints in attributes declared later have priority over constraints in attributes declared earlier
- `f3.s1_2`
- `f3.s1_1`
- `f2.s1_2`
- `f2.s1_1`
- `f1.s2_1`
- `f1.s1_2`
- `f1.s1_1`

Copyright © 2026 Accellera. All rights reserved.

When an attribute containing soft constraints is randomized via procedural randomization, only the soft constraints within the attribute being randomized and the inline constraints are considered from a priority perspective.

```
action Entry {
  rand S1 f1;
  rand S2 f2;

  exec post_solve {
    randomize f2 with {
        soft x in [10, 20, 30]; // w_1
        soft x < f1.x; // w_2
    }
  }
}
```

Example 161—Soft constraints and procedural randomization

In Example 161, only the soft constraints within *S2* (the type of attribute f2) and the in-line constraints are considered when randomizing *f2*. In this case, the relative priority (highest to lowest) of the constraints is as follows:

— `w_2`: Inline constraints have the highest relative priority
— `w_1`: Later-defined constraints have priority over earlier-defined constraints
— `f2.s2_1`
— `f2.s1_2`
— `f2.s1_1`

Soft constraints may be applied to aggregate data. An aggregate soft constraint is equivalent to individual soft constraints on each aggregate-attribute member.

```
struct S1 {
  rand bit[16] a, b;
}
action Entry {
  rand S1 f1;
  constraint soft f1 == {.a=0, .b=0}; // c1
  constraint soft f1.a == 0 && f1.b == 0; // c2
  constraint soft f1.a == 1; // c3
}
```

Example 162—Aggregate soft constraints

In Example 162, the aggregate constraint *c1* is equivalent to the two-subexpressions constraint *c2*. The higher-priority constraint *c3* causes the lower-priority soft constraints involving `a` to be discarded. This results in both sub-expressions of the constraint being discarded.

The following apply to how a PSS processing tool shall evaluate contradictions:

— In the presence of a constraint contradiction involving soft constraints, the PSS tool shall disregard soft constraints in priority order until either the contradiction is resolved or the set of soft constraints is exhausted.
— A contradiction can never result from a constraint set consisting entirely of soft constraints.

Copyright © 2026 Accellera. All rights reserved.

— A description involving soft constraints where no constraints contradict is equivalent to a description involving the same constraint statements as hard constraints.

13.1.13 Distribution directive

The distribution directive provides a value-distribution specification for a given expression to the constraint solver within the PSS processing tool. A distribution directive is a type of soft constraint. It has a lower priority than all other soft constraints.

```
struct S1 {
  rand bit[4] a1;
  constraint soft a1 == 5;
}
struct S2 {
  rand bit[4] a2;
  constraint soft a2 == 7;
}
struct S3 {
  rand S1 v1;
  rand S2 v2;
  constraint v1.a1 == v2.a2;
}
```

Example 163—Distribution directive soft constraint

In Example 163, the soft constraint in S2 has a higher priority than the soft constraint in S1. Consequently, v1.a1 and v2.a2 have the value 7.

```
constraint_body_item ::=
    ...
  | dist_directive
    ...
dist_directive ::= dist expression in [ dist_list ] ;
dist_list ::= dist_item { , dist_item }
dist_item ::= open_range_value [ dist_weight ]
dist_weight ::=
    := expression
  | :/ expression
```

Syntax 67—Distribution directive

A **dist** directive is a standalone statement from a syntax perspective. It is used to influence the value distribution of the target expression, but is not itself an expression.

The *dist_list* is a comma-separated list of integral expressions and ranges. Each term in the list can be given a non-negative weight, specified via the **:=** or **:/** operators. If no weight is specified for a given item, the default weight is **:= 1**.

In the absence of conflicting constraints, the value of the distribution target expression shall fall within the *dist_list*; the probability that the distribution target expression matches any value in the *dist_list* is

Copyright © 2026 Accellera. All rights reserved.

proportional to its specified weight. Constraints take priority over the **dist** directive and may force the distribution target expression to fall outside the set of values captured by the *dist_list*.

Value-distribution probability is only specified with respect to a single **dist** directive acting on an expression. In the presence of multiple **dist** directives acting on common expression elements with different distribution weights, the resulting value distribution across the common expression elements is undefined.

The **:=** operator assigns the specified weight to the item in the case of a single-value *dist_item*. In the case of a value-range *dist_item*, the weight is assigned to each value in the value range.

The **:/** operator assigns the specified weight to the item in the case of a single-value *dist_item*. In the case of a value-range *dist_item*, the weight is distributed across the values in the range. In other words, if there are *n* values in the range, each value will have a weight of *weight / n*.

The following also apply:

a) The left-hand expression shall be a scalar expression and contain at least one **rand** variable. Specifying a range shall only be allowed with a numeric or enumeration type.

b) **rand** variables may not be used in **dist** weights or value ranges.

c) The total weight associated with a value is the sum of all weights applied to that value in the *dist_list* using the **:=** and **:/** operators.

d) When the left-hand side expression and expressions and range bounds in *dist_list* are integers, their types shall be treated similarly to the set membership operator with a range-list on the right-hand side (see 8.7.1and Table 22).

13.1.13.1 Examples

```
struct S {
  rand bit[32] x;
  constraint dist x in [100..102 := 1, 200 := 2, 300 := 5];
}
```

Example 164—Distribution directive on single variable

In the example above, `x` is weighted to have a value range [100..102, 200, 300]. Additionally, value selection is weighted 1, 1, 1, 2, 5.

```
struct S {
  rand bit[32] x;
  constraint dist (x+6) in [100..102 := 1, 200 := 2, 300 := 5];
}
```

Example 165—Distribution directive on expression

Distribution weights may be applied to expressions as well as to individual variables. In the example above, the expression `(x+6)` is weighted to have a value range [100..102, 200, 300] with weights 1, 1, 1, 2, 5. Note that this is equivalent to applying the value ranges [94..96, 194, 294] to `x`.

Copyright © 2026 Accellera. All rights reserved.

```
struct S {
  rand bit[32] x;
  constraint dist x in [100..102 :/ 1, 200 := 2, 300 := 5];
}
```

Example 166—Distribution directive weight specification forms

In the example above, x is weighted to have a value range [100..102, 200, 300]. Additionally, value selection is weighted 1/3, 1/3, 1/3, 2, 5.

```
struct S {
  rand bit[32] x;
  bit          y;
  constraint dist x in [100..102 := 1, 200 := 2, 300 := 5];
  constraint (y==1) -> x > 300;
}
```

Example 167—Constraint priority over distribution directive

In the example above, a constraint can cause the value of x to be outside the **dist** directive range in some cases. When y is set to 1, the implication constraint prevents the **dist** directive from biasing the distribution of the target expression. This case does not result in a solve failure.

```
struct S {
  rand bit[32] x;
  bit[32]      w; // default value is 0
  constraint dist x in [100..102 := 1, 200 := 2, 300 := w];
}
```

Example 168—Zero-valued distribution weight

In Example 168 above, x is constrained using the declared value range of [100..102, 200, 300]. However, value 300 is given a weight of 0. Consequently, the effective value range of x will be [100..102, 200]. The value selection is weighted 1, 1, 1, 2.

13.2 Scheduling constraints

Scheduling constraints relate two or more actions or sub-activities from a scheduling point of view. Scheduling constraints do not themselves introduce new action traversals. Rather, they affect actions explicitly traversed in contexts that do not already dictate specific relative scheduling. Such contexts necessarily involve actions directly or indirectly under a **schedule** statement (see 11.3.5). Similarly, scheduling constraints can be applied to named sub-activities, see Syntax 68.

13.2.1 Syntax

```
activity_scheduling_constraint ::= constraint ( parallel | sequence )
  { hierarchical_id , hierarchical_id { , hierarchical_id } } ;
```

Syntax 68—Scheduling constraint statement

Copyright © 2026 Accellera. All rights reserved.

The following also apply:

a) **constraint sequence** schedules the related actions so that each completes before the next one starts (equivalent to a sequential activity block, see 11.3.3).

b) **constraint parallel** schedules the related actions such that they are invoked in a synchronized way and then proceed without further synchronization until their completion (equivalent to a parallel activity statement, see 11.3.4).

c) Scheduling constraints may not be applied to action handles that are traversed multiple times. In particular, they may not be applied to actions traversed inside an iterative statement: **repeat**, **repeat-while**, and **foreach** (see 11.4). However, the iterative statement itself, as a named sub-activity, can be related in scheduling constraints.

d) Scheduling constraints involving action-handle variables that are not traversed at all, or are traversed in branches not actually chosen from **select** or **if** statements (see 11.4), hold vacuously.

e) Scheduling constraints shall not undo or conflict with any scheduling requirements of the related actions.

13.2.2 Example

Example 169 demonstrates the use of a scheduling constraint. In it, compound action `my_sub_flow` specifies an activity in which action a is executed, followed by the group `b`, `c`, and `d`, with an unspecified scheduling relation between them. Action `my_top_flow` schedules two executions of `my_sub_flow`, relating their sub-actions using scheduling constraints.

```
action my_sub_flow {
   A a; B b; C c; D d;

   activity {
      sequence {
         a;
         schedule {
            b; c; d;
         };
      };
   };
};

action my_top_flow {
   my_sub_flow sf1, sf2;

   activity {
      schedule {
         sf1;
         sf2;
      };
   };

   constraint sequence {sf1.a, sf2.b};
   constraint parallel {sf1.b, sf2.b, sf2.d};
};
```

Example 169—Scheduling constraints

Copyright © 2026 Accellera. All rights reserved.

13.3 Sequencing constraints on state objects

A pool of **state** type stores exactly one state object at any given time during the execution of a test scenario, thus serving as a state variable (see 12.5). Any **action** that outputs a state object to a pool is considered a state transition with respect to that state variable. Within the context of a state type, reference can be made to attributes of the previous state, relating them in Boolean expressions to attributes values of this state. This is done by using the built-in reference variable **`prev`** (see 9.3.3).

NOTE—Any constraint in which **`prev`** occurs is vacuously satisfied in the context of the initial state object.

In Example 170, the first constraint in `power_state_s` determines that the value of `domain_B` may only decrement by `1`, remain the same, or increment by `1` between consecutive states. The second constraint determines that if a `domain_C` in any given state is `0`, the subsequent state has a `domain_C` of `0` or `1` and `domain_B` is `1`. These rules apply equally to the output of the two actions declared under component `power_ctrl_c`.

```
state power_state_s {
  rand int in [0..3] domain_A, domain_B, domain_C;

  constraint domain_B in { prev.domain_B - 1,
                           prev.domain_B,
                           prev.domain_B + 1};

  constraint prev.domain_C==0 -> domain_C in [0,1] || domain_B==0;
};
...
component power_ctrl_c {
  pool power_state_s psvar;
  bind psvar *;

  action power_trans1 {
    output power_state_s next_state;
  };

  action power_trans2 {
    output power_state_s next_state;
    constraint next_state.domain_C == 0;
  };
};
...
```

Example 170—Sequencing constraints

13.4 Randomization process

PSS supports randomization of plain-data type fields associated with scenario elements, as well as randomization of different relations between scenario elements, such as scheduling, resource allocation, and data flow. Moreover, the language supports specifying the order of random value selection, coupled with the flow of execution, in a compound action's sub-activity, the **activity** clause. Activity-based random value selection is performed with specific rules to simplify activity composition and reuse and minimize complexity for the user.

Random attribute fields of **struct** type are randomized as a unit. Traversal of a sub-action field triggers randomization of random attribute fields of the **action** and the resolution of its flow/resource object references. This is followed by evaluation of the action's activity if the action is compound.

Copyright © 2026 Accellera. All rights reserved.

13.4.1 Random attribute fields

This section describes the rules that govern whether an element is considered randomizable.

13.4.1.1 Semantics

a) Struct attribute fields qualified with the **rand** keyword are randomized if a field of that struct type is also qualified with the **rand** keyword.

b) Action attribute fields qualified with the **rand** keyword are randomized at the beginning of action execution. In the case of compound actions, **rand** attribute fields are randomized prior to the execution of the activity and, in all cases, prior to the execution of the action's *exec block*s (except **pre_solve**, see 13.4.12).

NOTE—It is often helpful to directly traverse attribute fields within an activity. This is equivalent to creating an intermediate action with a random attribute field of the plain-data type.

13.4.1.2 Examples

In Example 171, struct `S1` contains two attribute fields. Attribute field `a` is qualified with the **rand** keyword, while `b` is not. Struct `S2` creates two attribute fields of type `S1`. Attribute field `s1_1` is also qualified with the **rand** keyword. `s1_1.a` will be randomized, while `s1_1.b` will not. Attribute field `s1_2` is not qualified with the **rand** keyword, so neither `s1_2.a` nor `s1_2.b` will be randomized.

```
struct S1 {
  rand bit[4]   a;
  bit[4]        b;
}

struct S2 {
  rand S1         s1_1;
  S1              s1_2;
}
```

Example 171—Struct rand and non-rand fields

Example 172 shows two **action**s, each containing a **rand**-qualified data field (`A::a` and `B::b`). Action `B` also contains two fields of action type `A` (`a_1` and `a_2`). When action `B` is executed, a value is assigned to the random attribute field `b`. Next, the **activity** body is executed. This involves assigning a value to `a_1.a` and subsequently to `a_2.a`.

```
action A {
  rand bit[4]   a;
}

action B {
  A  a_1, a_2;
  rand bit[4]   b;

  activity {
    a_1;
    a_2;
  }
}
```

Example 172—Action rand-qualified fields

Copyright © 2026 Accellera. All rights reserved.

Example 173 shows an action-qualified field in action `B` named `a_bit`. The PSS processing tool assigns a value to `a_bit` when it is traversed in the `activity` body. The semantics are identical to assigning a value to the **rand**-qualified action field `A::a`.

```
action A {
    rand bit[4]   a;
  }

  action B {
    action bit[4] a_bit;
    A               a_1;

    activity {
      a_bit;
      a_1;
    }
  }
```

Example 173—Action-qualified fields

13.4.2 Randomization of lists

When a **rand**-qualified list variable is randomized, its elements are randomized and given values consistent with any constraints on them. The **`size`** of the array is not randomized, and may not be constrained (see 7.9.3.4).

Hierarchical constraint references to list elements can be declared in locations where it is not yet known whether the list element exists. Example 174 illustrates such a case.

```
action sub_a {
   rand list<bit[8]> lst;
   exec pre_solve {
      lst.push_back(0);
   }
}

action parent_a {
   sub_a a;
   rand int yy;
   constraint a.lst[0] == yy;

   activity {
     a;
   }
}
```

Example 174—Hierarchical constraint reference to list element

Constraints on list elements shall hold when the list is randomized. In this example, the list is randomized as part of the traversal of action handle `a`. At this point in time, the list contains a single element, and the constraint on this element is valid. If the referenced list element does not exist at the point of list randomization, then the PSS processing tool shall flag an error.

Copyright © 2026 Accellera. All rights reserved.

13.4.3 Randomization of flow objects

When an **action** is randomized, its **input** and **output** fields are assigned a reference to a flow object of the respective type. On entry to any of the action's *exec block*s (except **pre_solve**, see 20.1.2), as well as its **activity** clause(s), values for all **rand** data attributes accessible through its inputs and outputs fields are resolved. The values accessible in these contexts satisfy all constraints. Constraints can be placed on attribute fields from the immediate type context, from a containing struct or action at any level or via the input/output fields of actions.

The same flow object may be referenced by an action outputting it and one or more actions inputting it. The binding of inputs to outputs may be explicitly specified in an **activity** clause or may be left unspecified. In cases where binding is left unspecified, the counterpart action of a flow object's input/output may already be one explicitly traversed in an activity or it may be introduced implicitly by the PSS processing tool to satisfy the binding rules (see Clause 14). In the case where multiple actions input the same buffer object type, the input references may be constrained to indicate that they refer to the same object. In all of these cases, value selection for the data attributes of a flow object shall satisfy all constraints coming from the action that outputs it and actions that input it.

Consider the model in Example 175. Assume a scenario is generated starting from action `test`. The traversal of action `write1` is scheduled, followed by the traversal of action `read`. When `read` is randomized, its input `in_obj` shall be resolved. Every buffer object shall be the output of some action. The activity does not explicitly specify the binding of `read`'s input to any action's output, but it shall be resolved regardless. Action `write1` outputs a `mem_obj` whose `dat` is in the range `1` to `5`, due to a constraint in action `write1`. But, `dat` of the `mem_obj` instance `read` inputs shall be in the range `8` to `12`. So `read.in_obj` cannot be bound to `write1.out_obj` without violating a constraint. The PSS processing tool shall schedule another action of type `write2` at some point prior to `read`, whose `mem_obj` is bound to `read`'s input. In selecting the value of `read.in_obj.dat`, the PSS processing tool shall consider the following:

- `dat` is an even integer, due to the constraint in `mem_obj`.
- `dat` is in the range `6` to `10`, due to a constraint in `write2`.
- `dat` is in the range `8` to `12`, due to a constraint in `read`.

This restricts the legal values of `read.in_obj.dat` to either `8` or `10`.

Copyright © 2026 Accellera. All rights reserved.

```
component top {
    buffer mem_obj {
    rand int dat;
    constraint dat%2 == 0; // dat shall be even
    }

    action write1 {
        output mem_obj out_obj;
        constraint out_obj.dat in [1..5];
    }

    action write2 {
        output mem_obj out_obj;
        constraint out_obj.dat in [6..10];
    }

    action read {
        input mem_obj in_obj;
        constraint in_obj.dat in [8..12];
    }

    action test {
        activity {
            do write1;
            do read;
        }
    }
}
```

Example 175—Randomizing flow object attributes

13.4.4 Randomization of resource objects

When an **action** is randomized, its resource claim fields (of **resource** type declared with **lock** / **share** modifiers, see 9.4.1) are assigned a reference to a resource object of the respective type. On entry to any of the action's *exec block*s (except **pre_solve**, see 20.1.2) or its **activity** clause, values for all random attribute fields accessible through its resource fields are resolved. The same resource object may be referenced by any number of actions, given that no two concurrent actions lock it (see 9.4.2). Value selection for random attribute fields of a resource object satisfy constraints coming from all actions to which it was assigned, either in **lock** or **share** mode.

Consider the model in Example 176. Assume a scenario is generated starting from action `test`. In this scenario, three actions are scheduled to execute in parallel: `a1`, `a2`, and `a3`, followed sequentially by a traversal of `a4`. In the **parallel** statement, action `a3` of type `do_something_else` shall be exclusively assigned one of the two instances of resource type `rsrc_obj`, since `do_something_else` claims it in **lock** mode. Therefore, the other two actions, of type `do_something`, necessarily share the other instance. When selecting the value of attribute `kind` for that instance, the PSS processing tool considers the following constraints:

- `kind` is an enumeration whose domain has the values `A`, `B`, `C`, and `D`.
- `kind` is not `A`, due to a constraint in `do_something`.
- `a1.my_rsrc_inst` is referencing the same `rsrc_obj` instance as `a2.my_rsrc_inst`, as there would be a resource conflict otherwise between one of these actions and `a3`.
- `kind` is not `B`, due to an in-line constraint on `a1`.
- `kind` is not `C`, due to an in-line constraint on `a2`.

Copyright © 2026 Accellera. All rights reserved.

`D` is the only legal value for `a1.my_rsrc_inst.kind` and `a2.my_rsrc_inst.kind`.

Since there are only two instances of `rsrc_obj` in `rsrc_pool`, and one of the instances is claimed via the **share** in `a1` and `a2`, the other instance will be locked by `a3`. In order to determine the value of its `kind` field, we shall consider the in-line constraint on the traversal of `a4`. Since `a4.my_rsrc_inst.kind` is constrained to the value `A`, this shall be a different instance from the one shared by `a1` and `a2`. Therefore, this is the same instance that is claimed by `a3`, and therefore `a3.my_rsrc_inst.kind` shall also have the value of `A`.

```
component top {
    enum rsrc_kind_e {A, B, C, D};

    resource rsrc_obj {
        rand rsrc_kind_e kind;
    }

    pool[2] rsrc_obj rsrc_pool;
    bind rsrc_pool *;

    action do_something {
        share rsrc_obj my_rsrc_inst;
        constraint my_rsrc_inst.kind != A;
    }

    action do_something_else {
        lock rsrc_obj my_rsrc_inst;
    }

    action test {
        do_something      a1, a2;
        do_something_else a3, a4;
        activity {
            parallel {
                a1 with { my_rsrc_inst.kind != B; };
                a2 with { my_rsrc_inst.kind != C; };
                a3;
            }
            a4 with { my_rsrc_inst.kind == A; };
        }
    }
}
```

Example 176—Randomizing resource object attributes

13.4.5 Randomization of component assignment

When an **action** is randomized, its association with a component instance is determined. The built-in field **`comp`** is assigned a reference to the selected component instance. The assignment shall satisfy constraints where **`comp`** fields occur (see 9.1.5). Furthermore, the assignment of an action's **`comp`** field corresponds to the pools in which its inputs, outputs, and resources reside. If action `a` is assigned resource instance `r`, `r` is taken out the pool bound to `a`'s resource reference field in the context of the component instance assigned to `a`. If action `a` outputs a flow object which action `b` inputs, both output and input reference fields shall be bound to the same pool under `a`'s component and `b`'s component respectively. See Clause 12 for more on pool binding.

Copyright © 2026 Accellera. All rights reserved.

13.4.6 Procedural randomization of data

Procedural constrained randomization is performed using the **randomize** statement shown in Syntax 111.

The randomization target is composed of one or more variables of plain-data type. The entire set of variables is randomized together. Additional constraints may be added via the optional **with** block.

The set of variables and constraints involved in a procedural randomization statement is determined from the variables and in-line constraints passed to the statement. The variables and constraints described below are solved together.

- Randomization target variables are those that are specified as operands of the **randomize** statement. Target variables are treated as random, independent of whether they are declared **rand**.
- If a target variable is of a **struct** type, sub-fields declared **rand** are treated as random. Those not declared **rand** are treated as invariants.
- Constraints declared inside the target-variable types are applied.
- In-line constraints are applied.

```
struct S1 {
  rand bit[8] a, b;
}

struct S2 {
  rand S1 f1;
  S1 f2;
  constraint f1.a < f2.a;
}

action A {
  exec post_solve {
    S2      v1;
    bit[4]  v2;

    v1.f2.a = 100;
    randomize v1, v2 with {v1.f1.a < v2;}
  }
}
```

Example 177—Procedural randomization

In Example 177 above, `A::post_solve` performs procedural randomization on two variables (`v1`, `v2`):

a) `v1` is of **struct** type `S2`, and has two struct-type fields of the same type `S1`.
   1) `f1` is declared random.
   2) `f2` is declared non-random.
   3) A constraint is placed between sub-fields of `f1` and `f2`.

b) `v2` is of **bit**`[4]` type and thus has a maximum value of `15`.

An in-line constraint is placed between `v1.f1.a` and `v2`. When the procedural randomization statement executes, it considers:

a) Random variables: `v1.f1.a`, `v1.f1.b`, `v2`
b) Invariants: `v1.f2.a`, `v1.f2.b`
c) Invariant values:

Copyright © 2026 Accellera. All rights reserved.

1) `v1.f2.a == 100`
2) `v1.f2.b == 0`

d) Constraints:

1) `v1.f1.a < v2`
2) `v1.f1.a < v1.f2.a`

`v1.f1.a` will have a value `[0..14]` because it is required to be less than `100` (`v1.f2.a`) and less than the maximum value of `v2` (`15`).

13.4.6.1 Support on solve and target platforms

Support for procedural randomization in target **exec** blocks is restricted to built-in functions (e.g., **`urandom()`**) and randomization of scalar integer quantities. Randomization of **struct** data types is restricted to the solve platform, and may not be performed directly or indirectly from target **exec** blocks.

When procedural randomization is performed on the solve platform, any solve-time **exec** blocks within the scope of variables that are part of a procedural randomization are evaluated as part of the randomization process.

```
import std_pkg::*;

struct S1 {
  rand bit[8] a, b;
  exec pre_solve { print("Pre S1"); }
  exec post_solve { print("Post S1"); }
}

struct S2 {
  rand S1 f1;
  S1 f2;
  constraint f1.a < f2.a;
  exec pre_solve { print("Pre S2"); }
  exec post_solve { print("Post S2"); }
}

action A {
  exec post_solve {
    S2      v1;
    bit[4]  v2;

    v1.f2.a = 100;
    randomize v1, v2 with {v1.f1.a < v2;}
  }
}
```

Example 178—Evaluation of solve-time exec blocks in procedural randomization

In Example 178 above, we would expect to see the following when procedural randomization is invoked:

```
Pre S2
Pre S1
Pre S1
Post S2
Post S1
```

Copyright © 2026 Accellera. All rights reserved.

```
Post S1
```

13.4.6.2 Random stability

When procedural randomization features are used in *solve-time* **exec** blocks (**pre_solve**, **post_solve**, **pre_body**), random stability shall be ensured when the PSS description and the random seed specified to the PSS processing tool remain the same.

When procedural randomization features are used in *target* **exec** blocks (**body**), random stability shall be ensured when the PSS description, random seed specified to the runtime environment (if applicable), and design behavior remain the same.

13.4.7 Random value selection order

A PSS processing tool conceptually assigns values to sub-action fields of the **action** in the order they are encountered in the **activity**. On entry into an activity, the value of plain-data fields qualified with **action** and **rand** sub-fields of action-type fields are considered to be undefined.

Example 179 shows a simple activity with three action-type fields (`a`, `b`, `c`). A PSS processing tool might assign `a.val=2`, `b.val=4`, and `c.val=7` on a given execution.

```
action A {
  rand bit[4] val;
}

action my_action {
  A a, b, c;

  constraint abc_c {
    a.val < b.val;
    b.val < c.val;
  }
   activity {
    a;
    b;
    c;
  }
}
```

Example 179—Activity with random fields

13.4.8 Evaluation of expressions with action handles

Upon entry to an activity, all action handles (fields of action type) are considered uninitialized. Additionally, action handles previously traversed in an activity are reset to their uninitialized state upon entry to an activity block in which they are traversed again (an action handle may be traversed only once in any given activity scope and its nested scopes (see 11.3.1.1)). This applies equally to traversals of an action handle in a loop and to multiple occurrences of the same action handle in different activity blocks.

The value of all attributes reachable through uninitialized action handles, including direct attributes of the sub-actions and attributes of objects referenced by them, are unresolved. Only when all action handles in an expression are initialized, and all accessed attributes assume definite value, can the expression be evaluated.

Copyright © 2026 Accellera. All rights reserved.

Constraints accessing attributes through action handles are never violated. However, they are considered vacuously satisfied so long as these action handles are uninitialized. The Boolean expressions only need to evaluate to *true* at the point(s) in an activity when all action handles used in a constraint have been traversed.

Expressions in activity statements accessing attributes through action handles shall be illegal if they are evaluated at a point in which any of the action handles are uninitialized. Similarly, expressions in solve-exec (**pre_solve** and **post_solve**) statements of compound actions accessing attributes of sub-actions shall be illegal, since these are evaluated prior to the activity (see 13.4.12), and all action handles are uninitialized at that point. This applies equally to right-value and left-value expressions.

Example 180 shows a root action (`my_action`) with sub-action fields and an **activity** containing a loop. A value for `a.x` is selected, then two sets of values for `b.x` and `c.x` are selected.

```
action A {
  rand bit[4] x;
}

action my_action {
  A a, b, c;
  constraint abc_c {
    a.x < b.x;
    b.x < c.x;
  }
  activity {
    a;
    repeat (2) {
      b;
      c; // at this point constraint 'abc_c' shall hold non-vacuously
    }
  }
}
```

Example 180—Value selection of multiple traversals

The following breakout shows valid values that could be selected here:

Repetition	a.x	b.x	c.x
1	3	5	6
2	3	9	13

Note that `b.x` of the second iteration does not have to be less than `c.x` of the first iteration since action handle `c` is uninitialized on entry to the second iteration. Note also that similar behavior would be observed if the **repeat** would be unrolled, i.e., if the activity contained instead two blocks of `b`, `c` in sequence.

Example 181 demonstrates two cases of illegal access of action-handle attributes. In these cases, accessing sub-action attributes through uninitialized action handles shall be flagged as errors.

Copyright © 2026 Accellera. All rights reserved.

```
action A {
  rand bit[4] x;
  int y;
}

action my_action {
  A a, b, c;

  exec post_solve {
    a.y = b.x; // ERROR - cannot access uninitialized action handle
   attributes
  }

  activity {
    a;
    if (a.x > 0) { // OK - 'a' is resolved
      b;
      c;
    }
    {
      if (c.y == a.x) { // ERROR - cannot access attributes of
                        // uninitialized action handle 'c.y'
        b;
      }
      c;
    }
  }
}
```

Example 181—Illegal accesses to sub-action attributes

13.4.9 Relationship lookahead

Values for random fields in an **activity** are selected and assigned as the fields are traversed. When selecting a value for a random field, a PSS processing tool shall take into account both the explicit constraints on the field and the implied constraints introduced by constraints on those fields traversed during the remainder of the activity traversal (including those introduced by inferred actions, binding, and scheduling). This rule is illustrated by Example 182.

13.4.9.1 Examples

Example 182 shows a simple **struct** with three random attribute fields and constraints between the fields. When an instance of this struct is randomized, values for all the random attribute fields are selected at the same time.

```
struct abc_s {
   rand bit[4] in [0..12] a_val, b_val, c_val;

   constraint {
      a_val < b_val;
      b_val < c_val;
   }
}
```

Example 182—Struct with random fields

Copyright © 2026 Accellera. All rights reserved.

Example 183 shows a root action (`my_action`) with three sub-action fields and an activity that traverses these sub-action fields. It is important that the random-value selection behavior of this activity and the **struct** shown in Example 182 are the same. If a value for `a.val` is selected without knowing the relationship between `a.val` and `b.val`, the tool could select `a.val=15`. When `a.val=15`, there is no legal value for `b.val`, since `b.val` shall be greater than `a.val`.

a) When selecting a value for `a.val`, a PSS processing tool shall consider the following:
   1) `a.val` is in the range `0` to `15`, due to its domain.
   2) `b.val` is in the range `0` to `15`, due to its domain.
   3) `c.val` is in the range `0` to `15`, due to its domain.
   4) `a.val < b.val`.
   5) `b.val < c.val`.

   This restricts the legal values of `a.val` to `0` to `13`.

b) When selecting a value for `b.val`, a PSS processing tool shall consider the following:
   1) The value selected for `a.val`.
   2) `b.val` is in the range `0` to `15`, due to its domain.
   3) `c.val` is in the range `0` to `15` due to its domain.
   4) `a.val < b.val`.
   5) `b.val < c.val`.

```
action A {
  rand bit[4] val;
}

action my_action {
  A a, b, c;

  constraint abc_c {
    a.val < b.val;
    b.val < c.val;
  }
   activity {
    a;
    b;
    c;
  }
}
```

Example 183—Activity with random fields

13.4.10 Lookahead and sub-actions

Lookahead shall be performed across traversal of sub-action fields and shall comprehend the relationships between action attribute fields.

Example 184 shows an action named `sub` that has three sub-action fields of type `A`, with constraint relationships between those field values. A top-level action has a sub-action field of type `A` and type `sub`, with a constraint between these two action-type fields. When selecting a value for the `top_action.v.val` random attribute field, a PSS processing tool shall consider the following:

- `top_action.s1.a.val == top_action.v.val`
- `top_action.s1.a.val < top_action.s1.b.val`

Copyright © 2026 Accellera. All rights reserved.

This implies that `top.v.val` shall be less than `14` to satisfy the `top_action.s1.a.val < top_action.s1.b.val` constraint.

```
component top {
    action A {
        rand bit[4] val;
    }

    action sub {
        A a, b, c;

        constraint abc_c {
            a.val < b.val;
            b.val < c.val;
        }

        activity {
            a;
            b;
            c;
        }
    }

    action top_action {
        A v;
        sub s1;

        constraint c {
            s1.a.val == v.val;
        }

        activity {
            v;
            s1;
        }
    }
}
```

Example 184—Sub-activity traversal

13.4.11 Lookahead and generic constraints

Generic constraints introduce traversal-dependent constraints. A PSS processing tool shall account for these additional constraints when making random attribute field value selections. A generic constraint shall hold for the entire activity branch on which it is referenced, as well as to the remainder of the activity.

Example 185 shows an activity with two generic constraints that are mutually exclusive. If the first branch is selected, `b.val <= 5` and `b.val < a.val`. If the second branch is selected, `b.val <= 7` and `b.val > a.val`. A PSS processing tool shall select a value for `a.val` such that a legal value for `b.val` also exists (presuming this is possible).

Given the generic constraints, legal value ranges for `a.val` are `1` to `15` for the first branch and `0` to `6` for the second branch.

Copyright © 2026 Accellera. All rights reserved.

```
action A {
  rand bit[4] val;
}

action dyn {
  A        a, b;

  constraint d1 {
    b.val < a.val;
    b.val <= 5;
  }

  constraint d2 {
    b.val > a.val;
    b.val <= 7;
  }

  activity {
    a;
    select {
      d1;
      d2;
    }
    b;
  }
}
```

Example 185—Activity with generic constraints

13.4.12 Pre_solve and post_solve exec blocks

The **pre_solve** and **post_solve** *exec block*s enable external code to participate in the solve process. **pre_solve** and **post_solve** *exec block*s may appear in **struct** and **action** type declarations.

Statements in **pre_solve** blocks are used to set non-random attribute fields that are subsequently read by the solver during the solve process. Statements in **pre_solve** blocks can read the values of non-random attribute fields and their non-random children. Statements in **pre_solve** blocks cannot access handle-type fields (**input**/**output**, **lock**/**share**, action handles) or their children since these fields are null handles prior to the completion of randomization. Accessing plain-data random fields (e.g., **bit**, **int**, **struct**) is permitted. Reading the value of these fields in **pre_solve** blocks returns the initial value of the field. Values written to scalar plain-data random fields in **pre_solve** will be overwritten by the solve process.

Statements in **post_solve** blocks are evaluated after the solver has resolved values for random attribute fields and are used to set the values for non-random attribute fields based on randomly-selected values.

The execution order of **pre_solve** and **post_solve** *exec block*s, respectively, corresponds to the order random attribute fields are assigned by the solver. The ordering rules are as follows:

a) Order within a compound action is top-down—both the **pre_solve** and **post_solve** *exec block*s, respectively, of a containing action are executed before any of its sub-actions are traversed, and, hence, before the **pre_solve** and **post_solve**, respectively, of its sub-actions.

b) Order between actions follows their relative scheduling in the scenario: if action a_1 is scheduled before a_2, a_1's **pre_solve** and **post_solve** blocks, if any, are called before the corresponding block of a_2.

Copyright © 2026 Accellera. All rights reserved.

c) Order for flow objects (instances of struct types declared with a **buffer**, **stream**, or **state** modifier) follows the order of their flow in the scenario: a flow object's **pre_solve** or **post_solve** *exec block* is called after the corresponding *exec block* of its outputting action and before that of its inputting action(s).

d) A resource object's **pre_solve** or **post_solve** *exec block*s are called before the corresponding *exec block*(s) of all actions referencing it, regardless of their use mode (**lock** or **shared**).

e) Order within an aggregate data type (nested struct and collection fields) is top-down—the *exec block*s of the containing instance are executed before those of the contained.

PSS does not specify the execution order in other cases. In particular, any relative order of execution for sibling random **struct** attributes is legitimate and so is any order for actions scheduled in parallel where no flow objects are exchanged between them.

See 20.1 for more information on the *exec block* construct.

13.4.12.1 Examples

Example 186 shows a top-level struct `S2` that has rand and non-rand scalar fields, as well as two fields of struct type `S1`. When an instance of `S2` is randomized, the *exec block* of `S2` is evaluated first, but the execution for the two `S1` instances can be in any order. The following is one such possible order:

a) **pre_solve** in `S2`

b) **pre_solve** in `S2.s1_2`

c) **pre_solve** in `S2.s1_1`

d) assignment of attribute values

e) **post_solve** in `S2`

f) **post_solve** in `S2.s1_1`

g) **post_solve** in `S2.s1_2`

Copyright © 2026 Accellera. All rights reserved.

```
function bit[6] get_init_val();
function bit[6] get_exp_val(bit[6] stim_val);

struct S1 {
   bit[6] init_val;
   rand bit[6] rand_val;
   bit[6] exp_val;

   exec pre_solve {
      init_val = get_init_val();
   }

   constraint rand_val_c {
      rand_val <= init_val+10;
   }

   exec post_solve {
      exp_val = get_exp_val(rand_val);
   }
}

struct S2 {
   bit[6] init_val;
   rand bit[6] rand_val;
   bit[6] exp_val;

   rand S1 s1_1, s1_2;

   exec pre_solve {
      init_val = get_init_val();
   }

   constraint rand_val_c {
      rand_val > init_val;
   }

   exec post_solve {
      exp_val = get_exp_val(rand_val);
   }
}
```

Example 186—Pre_solve/post_solve

Example 187 illustrates the relative order of execution for **post_solve** *exec block*s of a containing action `test`, two sub-actions: `read` and `write`, and a buffer object exchanged between them.

The calls therein are executed as follows:

a) **post_solve** in `test`
b) **post_solve** in `write`
c) **post_solve** in `mem_obj`
d) **post_solve** in `read`

Copyright © 2026 Accellera. All rights reserved.

```
buffer mem_obj {
  exec post_solve { ... }
};

action write {
  output mem_obj out_obj;
  exec post_solve { ... }
};

action read {
  input mem_obj in_obj;
  exec post_solve { ... }
};

action test {
  write wr;
  read  rd;

  activity {
    wr;
    rd;
    bind wr.out_obj rd.in_obj;
  }
  exec post_solve { ... }
};
```

Example 187—Post_solve ordering between action and flow objects

13.4.13 Body blocks and sampling external data

exec body blocks, or functions invoked by them, can assign values to attribute fields. **exec body** blocks are evaluated for atomic actions as part of the test execution on the target platform (see 20.1). The impact of any field values modified by an **exec body** block is evaluated after the entire **exec body** block has completed.

Example 188 shows an **exec body** block that assigns two non-rand attribute fields. The impact of the new values applied to `y1` and `y2` are evaluated against the constraint system after the **exec body** block completes execution. It shall be illegal if the new values of `y1` and `y2` conflict with other attribute field values and constraints. Backtracking is not performed.

Copyright © 2026 Accellera. All rights reserved.

```
function bit[4] compute_val1(bit[4] v);
function bit[4] compute_val2(bit[4] v);
component pss_top {

  action A {
    rand bit[4] x;
    bit[4] y1, y2;

    constraint assume_y_c {
      y1 >= x && y1 <= x+2;
      y2 >= x && y2 <= x+3;

      y1 <= y2;
    }

    exec body {
      y1 = compute_val1(x);
      y2 = compute_val2(x);
    }
  }
}
```

Example 188—Exec body block sampling external data

Copyright © 2026 Accellera. All rights reserved.

14. Action inferencing

Perhaps the most powerful feature of PSS is the ability to focus purely on the user's verification intent, while delegating the means to achieve that intent. Previous clauses have introduced the semantic concepts to define such abstract specifications of intent. The modeling constructs and semantic rules thus defined for a portable stimulus model allow a tool to generate a number of scenarios from a single (partial) specification to implement the desired intent.

Beginning with a root action, which may contain an activity, a number of actions and their relative scheduling constraints is used to specify the verification intent for a given model. The other elements of the model, including flow objects, resources and their binding, as well as algebraic constraints throughout, define a set of rules that shall be followed to generate a valid scenario matching the specified intent. It is possible to fully specify a verification intent model, in which only a single valid scenario of actions may be generated. The randomization of data fields in the actions and their respective flow and resource objects would render this scenario as what is generally referred to as a "directed random" test, in which the actions are fully defined, but the data applied through the actions is randomized. The data values themselves may also be constrained so that there is only one scenario that may be generated, including fully-specified values for all data fields, in which case the scenario would be a "directed" test.

There are a number of ways to specify the scheduling relationship between actions in a portable stimulus model. The first, which allows explicit specification of verification intent, is via an activity. As discussed in Clause 11, an activity may define explicit scheduling dependencies between actions, which may include statements, such as **schedule**, **select**, **if-else** and others, to allow multiple scenarios to be generated even for a fully-specified intent model. Consider Example 189.

```
component pss_top {
  buffer data_buff_s {
    rand int val;
  };
  pool data_buff_s data_mem;
  bind data_mem *;

  action A_a {output data_buff_s dout;};
  action B_a {output data_buff_s dout;};
  action C_a {input  data_buff_s din;};
  action D_a {input  data_buff_s din;};

  action root_a {
    A_a a;
    B_a b;
    C_a c;
    D_a d;
    activity {
      select {a; b;}
      select {c; d;}
    }
  }
}
```

Example 189—Generating multiple scenarios

While an activity may be used to fully express the intent of a given model, it is more often used to define the critical actions that shall occur to meet the verification intent while leaving the details of how the actions may interact unspecified. In this case, the rules defined by the rest of the model, including flow object

Copyright © 2026 Accellera. All rights reserved.

requirements, resource limitations, pool bindings, and algebraic constraints, permit a tool to introduce the traversal of additional actions as defined by the model to ensure the generation of a valid scenario that meets the critical intent as defined by the activity. The introduction of an action in the execution of a scenario to complete a partially specified flow is called *action inferencing*.

The evaluation ordering rules for **pre_solve** and **post_solve** exec blocks of actions, objects, and structs, as specified in 13.4.12, apply regardless of whether the actions are explicitly traversed or inferred, and whether objects are explicitly or implicitly bound. In particular, the order conforms to the scheduling relations between **action**s, such that if an action is scheduled before another, its **pre_solve** and **post_solve** execs are evaluated before the other's. Backtracking is not performed across exec blocks. Assignments in **exec** blocks to attributes that figure in constraints may therefore lead to unsatisfied constraint errors. This applies to inferred parts of the scenarios in the same way as to parts that are explicitly specified in activities.

14.1 Implicit binding and action inferences

In a scenario description, the explicit binding of outputs to inputs may be left unspecified. In these cases, an implementation shall execute a scenario that reflects a valid completion of the given partial specification in a way that conforms to pool binding rules. If no valid scenario exists, the tool shall report an error. Completing a partial specification may involve decisions on output-to-input binding of flow objects in actions that are explicitly traversed. It may also involve introducing the traversal of additional actions, beyond those explicitly traversed, to serve as the counterpart of a flow object exchange. Once an action traversal is inferred to complete a given flow object exchange, it may also be considered for completing other flow object exchanges with which it may also be compatible.

Action inferences are necessary to make a scenario execution legal if the following conditions hold:

a) An input of any kind is not explicitly bound to an output, or an output of stream kind is not explicitly bound to an input.
b) There is no action explicitly traversed or inferred that is available to legally bind its output/input to the unbound input/output, i.e.,
    1) There is no action that is or may be scheduled before the inputting action in the case of buffer or state objects.
    2) There is no action that is or may be scheduled in parallel to the inputting/outputting action in the case of stream objects.

The inferencing of actions may be based on random or policy-driven (which may include specified coverage goals) decisions of a processing tool. Actions may only be inferred to complete a partially-specified flow. If all required input-to-output bindings are specified by explicit bindings to the traversed actions in the activity, an implementation may not introduce additional actions in the execution. See Annex E for more details on inference rules.

Consider the model in Example 190.

If action `send_data` is designated as the root action, this is clearly a case of partial scenario description, since action `send_data` has an input and an output, neither of which is explicitly bound. The buffer input `src_data` is bound to the `data_mem` object pool, so there shall be a corresponding output object also bound to the same pool to provide the buffer object. The only action type outputting an object of the required type that is bound to the same object pool is `load_data`. Thus, an implementation shall infer the prior traversal of `load_data` before traversing `send_data`.

Similarly, `load_data` has a state input that is bound to the `config_var` pool. Since the output objects of action types `setup_A` and `setup_B` are also bound to the same pool, `load_data.curr_cfg` can be bound to the output of either `setup_A` or `setup_B`, but cannot be the initial state due to the constraint in

Copyright © 2026 Accellera. All rights reserved.

load_data. In the absence of other constraints, the choice of whether to infer setup_A or setup_B may be randomized and the chosen action traversal shall occur before the traversal of load_data.

Moreover, send_data has a stream output out_data, which shall be bound to the corresponding input of another action that is also bound to the data_bus pool. So, an implementation shall infer the traversal of an action of type receive_data in parallel to send_data.

```
component pss_top {
  state config_s {};
  pool config_s config_var;
  bind config_var *;

  buffer data_buff_s {};
  pool data_buff_s data_mem;
  bind data_mem *;

  stream data_stream_s {};
  pool data_stream_s data_bus;
  bind data_bus *;

  action setup_A {
    output config_s new_cfg;
  };

  action setup_B {
    output config_s new_cfg;
  };

  action load_data {
    input config_s curr_cfg;
    constraint !curr_cfg.initial;
    output data_buff_s out_data;
  };

  action send_data {
    input data_buff_s src_data;
    output data_stream_s out_data;
  };

  action receive_data {
    input data_stream_s in_data;
  };
};
```

Example 190—Action inferences for partially-specified flows

Note that action inferences may be more than one level deep. The scenario executed by an implementation shall be the transitive closure of the specified scenario per the flow object dependency relations. Consider adding another action within the **pss_top** component in Example 190, e.g.,

```
action xfer_data {
  input data_buff_s src_data;
  output data_buff_s out_data;
};
```

Copyright © 2026 Accellera. All rights reserved.

In this case, the xfer_data action could also be inferred, along with setup_A or setup_B to provide the data_buff_s input to send_data.src_data. If xfer_data were inferred, then its src_data input would require the additional inference of another instance of setup_A, setup_B, or xfer_data to provide the data_buff_s. This "inference chain" would continue until either an instance of setup_A or setup_B is inferred, which would require no further inferencing, or the inferencing limit of the tool is exceeded, in which case an error would be reported.

Since the type of the inferred action is randomly selected from all available compatible action types, a tool may ensure that either setup_A or setup_B gets inferred before the inferencing limit is exceeded.

Consider Example 191. Starting with the constr_test action, two instances of the get_data action are traversed in parallel. Since each instance inputs a buffer of type data_buff_s, at least one instance of load_data shall be inferred to provide the input buffer. The equality constraint c2 requires that gd1.src_data and gd2.src_data are actually the same object, so only a single instance of load_data will be inferred. Without the c2 constraint, it would have been possible to infer two separate instances of load_data, each of which would provide a buffer object to either gd1 or gd2, although inferring a single instance is also legal. Note that the c1 constraint by itself is not sufficient to guarantee a single instance inference since there could be two distinct buffers with identical contents. With the c2 constraint present, the c1 constraint is redundant (but legal).

```
component pss_top {
  buffer data_buff_s {bit[4] val;};
  pool   data_buff_s dbuf_p;
  bind   dbuf_p *;

  action load_data {
    output data_buff_s out_data;
  }

  action get_data {
    input  data_buff_s src_data;
  }

  action constr_test {
    get_data gd1, gd2;

    constraint c1 {gd1.src_data.val == gd2.src_data.val;}
    constraint c2 {gd1.src_data     == gd2.src_data;}

    activity {
      parallel {gd1; gd2;}
    }
  }
}
```

Example 191—Buffer equality constraint to limit inferencing

Consider Example 192. In the constr_rsrc_test action, two instances of the m2m action are scheduled and traversed, each of which inputs and outputs a data_buff_s buffer object and locks a dma_descr resource object, followed by the parallel traversal of two instances of the get_data action. Constraint c3 ensures that both m2m instances input the same data_buff_s object and therefore a single instance of either load_data or m2m is inferred to provide it. Constraint c4 guarantees that the two get_data instances will each consume a different data_buff_s object, so each will be provided by either m2m1 or m2m2. Constraint c5 requires the two m2m instances to claim the same resource object, so the schedule

Copyright © 2026 Accellera. All rights reserved.

statement shall require one instance to be traversed before the other, in either order. Note that the commented-out constraint `c6` is equivalent to `c5`.

```
component pss_top {
  buffer    data_buff_s {bit[4] val;};
  resource  dma_descr   {bit[4] chan;};
  pool      data_buff_s dbuf_p;
  bind      dbuf_p *;
  pool [16] dma_descr   descr_p;
  bind      descr_p *;

  action load_data {
    output data_buff_s out_data;
  }

  action get_data {
    input  data_buff_s src_data;
  }

  action m2m {
    input  data_buff_s ibuf;
    output data_buff_s obuf;
    lock   dma_descr   descr;
  }

  action constr_rsrc_test {
    get_data gd1 , gd2;
    m2m      m2m1, m2m2;

    constraint c3 {m2m1.ibuf    == m2m2.ibuf;}
    constraint c4 {gd1.src_data != gd2.src_data;}
    constraint c5 {m2m1.descr   == m2m2.descr;}
//  constraint c6 {m2m1.descr.instance_id == m2m2.descr.instance_id;}

    activity {
      schedule {m2m1; m2m2;}
      parallel {gd1 ; gd2; }
    }
  }
}
```

Example 192—Resource equality constraint may affect scheduling

14.2 Object pools and action inferences

Action traversals may be inferred to support the flow object requirements of actions that are traversed in the model, whether they are explicitly traversed or inferred. The set of actions from which a traversal may be inferred is determined by object pool bindings.

In Example 193, there are two object pools of type `data_buff_s`, each of which is bound to a different set of object field references. The **select** statement in the activity of `root_a` will randomly choose either `c` or `d`, each of which has a `data_buff_s` buffer input type that requires a corresponding action to be inferred to supply the buffer object. Since `C_a` is bound to the same pool as `A_a`, if the generated scenario chooses `c`, then an instance of `A_a` shall be inferred to supply the `c.din` buffer input. Similarly, if `d` is chosen, then an instance of `B_a` shall be inferred to supply the `d.din` buffer input.

Copyright © 2026 Accellera. All rights reserved.

```
component pss_top {
  buffer data_buff_s {...};
  pool data_buff_s data_mem1, data_mem2;
  bind data_mem1 {A_a.dout, C_a.din};
  bind data_mem2 {B_a.dout, D_a.din};

  action A_a {output data_buff_s dout;};
  action B_a {output data_buff_s dout;};
  action C_a {input  data_buff_s din;};
  action D_a {input  data_buff_s din;};

  action root_a {
    C_a c;
    D_a d;
    activity {
      select {c; d;}
    }
  }
}
```

Example 193—Object pools affect inferencing

Consider the following modified version of Example 51 from 9.1.5.2. In this example, the traversal of action `foo` in the activity of action `gr_a` requires the inference of an action that can be bound to the same pool as `graphics::foo` and supply the compatible `bar_s` type flow object. Since the `bar_p` pool is bound by default to all components under `graphics` and `bus_c`, it is legal to infer the traversal of `bus_c::write` in parallel with `foo`, even though it was illegal to traverse this action explicitly as shown in Example 51.

```
component bus_c {
  import bar_pkg::*;
  action write{input bar_s b;...} // bar_s is a stream
}

component graphics {
  import bar_pkg::*;
  action foo {output bar_s b;...}
  action gr_a {
    activity {
        do foo; // will infer traversal of bus_c::write
                // to complete stream object connection
    }
  }
}

component pss_top {
  import bar_pkg::*;
  bus_c a0;
  graphics g;
  pool bar_s bar_p;
  bind bar_p *;
}
```

Example 194—Inferred traversal of an action outside of the containing component hierarchy

Copyright © 2026 Accellera. All rights reserved.

14.3 Data constraints and action inferences

As mentioned in Clause 13, introducing data constraints on flow objects or other elements of the design may affect the inferencing of actions. Consider a slightly modified version of Example 189, as shown in Example 195.

Since the explicit traversal of c does not constrain the `val` field of its input, it may be bound to the output of either explicitly traversed action `a` or `b`; thus, there are two legal scenarios to be generated with the second **select** statement evaluated to traverse action `c`. However, since the data constraint on the traversal of action `d` is incompatible with the in-line data constraints on the explicitly-traversed actions `a` or `b`, another instance of either `A_a` or `B_a` shall be inferred whose output shall be bound to `d.din`. Since there is no requirement for the buffer output of either `a` or `b` to be bound, one of these actions shall be traversed from the first **select** statement, but no other action shall be inferred.

```
component pss_top {
  buffer data_buff_s {
    rand int val;
  };
  pool data_buff_s data_mem;
  bind data_mem *;

  action A_a {output data_buff_s dout;};
  action B_a {output data_buff_s dout;};
  action C_a {input  data_buff_s din;};
  action D_a {input  data_buff_s din;};

  action root_a {
    A_a a;
    B_a b;
    C_a c;
    D_a d;
    activity {
      select {a with{dout.val<5;}; b with {dout.val<5;};}
      select {c; d with {din.val>5;};}
    }
  }
}
```

Example 195—In-line data constraints affect action inferencing

Copyright © 2026 Accellera. All rights reserved.

Consider, instead, if the in-line data constraints were declared in the action types, as shown in Example 196.

In this case, there is no valid action type available to provide the `d.din` input that satisfies its constraint as defined in the `D_a` action declaration, since the only actions that may provide the `data_buff_s` type, actions `A_a` and `B_a`, have constraints that contradict the input constraint in `D_a`. Therefore, the only legal action to traverse in the second select statement is `c`. In fact, it would be illegal to traverse action `D_a` under any circumstances for this model, given the contradictory data constraints on the flow objects.

```
component pss_top {
  buffer data_buff_s {
    rand int val;
  };
  pool data_buff_s data_mem;
  bind data_mem *;

  action A_a {
    output data_buff_s dout;
    constraint {dout.val<5;}
  };
  action B_a {
    output data_buff_s dout;
    constraint {dout.val<5;}
  };
  action C_a {
    input data_buff_s din;
  };
  action D_a {
    input data_buff_s din;
    constraint {din.val > 5;}
  };

  action root_a {
    A_a a;
    B_a b;
    C_a c;
    D_a d;
    activity {
      select {a; b;}
      select {c; d;}
    }
  }
}
```

Example 196—Data constraints affect action inferencing

Copyright © 2026 Accellera. All rights reserved.

15. Data coverage

The legal state space for all non-trivial verification problems is very large. Coverage goals identify important scenarios and key value ranges and value combinations that need to occur in order to exercise key functionality. Covering scenarios is the subject of the behavioral coverage, and it is described in Clause 16. Covering value ranges and combinations is called data coverage, and it is described in this clause. The **covergroup** construct is used to specify the data coverage targets.

The coverage targets specified by the **covergroup** construct are more directly related to the test scenario being created. As a consequence, in many cases the coverage targets would be considered coverage targets on the "generation" side of stimulus. PSS also allows data to be sampled by calling external functions. Coverage targets specified on data fields set by external functions can be related to the system state.

15.1 Defining the coverage model: covergroup

The **covergroup** construct encapsulates the specification of a coverage model. Each **covergroup** specification can include the following elements:

- A set of coverage points
- Cross coverage between coverage points
- Optional formal arguments
- Coverage options

The **covergroup** construct is a user-defined type. There are two forms of the **covergroup** construct. The first form allows an explicit type definition to be written once and instantiated multiple times in different contexts. The second form allows an in-line specification of an anonymous **covergroup** type and a single instance.

a) An *explicit* **covergroup** type can be defined in a **package**, **component**, **action**, **monitor**, or **struct**. In order to be reusable, an explicit **covergroup** type shall specify a list of formal parameters and shall not reference fields in the scope in which it is declared. An instance of an explicit **covergroup** type can be created in an **action**, **monitor**, or **struct**. Syntax 69 defines an explicit **covergroup** type.

b) An *in-line* **covergroup** can be defined in an **action**, **monitor**, or **struct** scope. An in-line covergroup can reference fields in the scope in which it is defined. 15.2 contains more information on in-line covergroups.

Copyright © 2026 Accellera. All rights reserved.

15.1.1 Syntax

The syntax for **covergroup**s is shown in Syntax 69.

```
covergroup_declaration ::=
    covergroup covergroup_identifier ( covergroup_port {, covergroup_port } )
    { {covergroup_body_item} }
covergroup_port ::= data_type identifier
covergroup_body_item ::=
    covergroup_option
  | covergroup_coverpoint
  | covergroup_cross
  | covergroup_body_compile_if
  | annotation_usage covergroup_body_item
  | stmt_terminator
covergroup_option ::=
    option . identifier = constant_expression ;
```

Syntax 69—Covergroup declaration

The following also apply:

a) The identifier associated with the **covergroup** declaration defines the name of the coverage model type.

b) A **covergroup** can contain one or more coverage points. A *coverage point* can cover a variable or an expression.

c) Each coverage point includes a set of bins associated with its sampled value. The bins can be user-defined or automatically created by a tool. Coverage points are detailed in 15.3.

d) A **covergroup** can specify cross coverage between two or more coverage points or variables. Any combination of more than two variables or previously declared coverage points is allowed. See also Example 198.

e) A **covergroup** can also specify one or more options to control how coverage data are structured and collected. Coverage options can be specified for the **covergroup** as a whole or for specific items within the **covergroup**, i.e., any of its coverage points or crosses. In general, a coverage option specified at the **covergroup** level applies to all of its items unless overridden in a specific item's definition. Coverage options are described in 15.5.

15.1.2 Examples

Example 197 defines an in-line covergroup `cs1` with a single coverage point labeled `c` associated with struct field `color`. The value of the variable `color` is sampled at the default sampling point: the end of an action's traversal in which the field `color` is randomized. Sampling is discussed in more detail in 15.6.

Because the coverage point does not explicitly define any bins, the tool automatically creates three bins, one for each possible value of the enumeration type. Automatic bins are described in 15.3.4.

Copyright © 2026 Accellera. All rights reserved.

```
enum color_e {red, green, blue};

struct s {
   rand color_e color;

   covergroup {
      c: coverpoint color;
   } cs1;
}
```

Example 197—Single coverage point

Example 198 creates an in-line covergroup `cs2` that includes two coverage points and two cross coverage items. Explicit coverage points labeled `Offset` and `Hue` are defined for variables `pixel_offset` and `pixel_hue`. PSS implicitly declares coverage points for variables `color` and `pixel_adr` to track their cross coverage. Implicitly declared coverage points are described in 15.4.

```
enum color_e {red, green, blue};

struct s {
    rand color_e        color;
    rand bit[4]       pixel_adr, pixel_offset, pixel_hue;

    covergroup {
        Hue : coverpoint pixel_hue;
        Offset : coverpoint pixel_offset;
        AxC: cross color, pixel_adr;
        all : cross color, Hue, Offset;
    } cs2;
}
```

Example 198—Two coverage points and cross coverage items

15.2 Covergroup instantiation

A **covergroup** type can be instantiated in **struct**, **action**, **monitor**, and **cover** contexts. If the **covergroup** declared formal parameters, these shall be bound to variables visible in the instantiation context. Instance-specific coverage options (see 15.5) may be specified as part of instantiation. If a **covergroup** is specific to the containing type, it cannot be generally instantiated in other types. In these cases, it is possible to declare a covergroup instance in-line. In this case, the **covergroup** type is *anonymous*.

Copyright © 2026 Accellera. All rights reserved.

15.2.1 Syntax

Syntax 70 specifies how a **covergroup** is instantiated and how an in-line covergroup instance is declared.

```
covergroup_instantiation ::=
    covergroup_type_instantiation
  | inline_covergroup
inline_covergroup ::= covergroup { { covergroup_body_item } } identifier ;
covergroup_type_instantiation ::= covergroup_type_identifier covergroup_identifier
    ( covergroup_portmap_list ) covergroup_options_or_empty
covergroup_type_identifier ::= type_identifier
covergroup_portmap_list ::=
    covergroup_portmap { , covergroup_portmap }
  | hierarchical_id_list
covergroup_portmap ::= . identifier ( hierarchical_id )
covergroup_options_or_empty ::=
    with { { covergroup_option } }
  | ;
```

Syntax 70—Instantiating a covergroup

15.2.2 Examples

Example 199 defines a covergroup type with a formal parameter list and creates a covergroup instance.

```
enum color_e {red, green, blue};

struct s {
    rand color_e color;

    covergroup cs1(color_e c) {
        c : coverpoint c;
    }

    cs1 cs1_inst(color);
}
```

Example 199—Creating and instantiating a covergroup type with a formal parameter list

Copyright © 2026 Accellera. All rights reserved.

Example 200 defines a covergroup type and creates a covergroup instance with instance-specific options.

```
enum color_e {red, green, blue};

struct s {
    rand color_e color;

    covergroup cs1 (color_e color) {
        c: coverpoint color;
    }

    cs1 cs1_inst (color) with {
        option.at_least = 2;
    };
}
```

Example 200—Creating a covergroup instance with instance-specific options

Example 201 creates an in-line covergroup instance.

```
enum color_e {red, green, blue};

struct s {
    rand color_e color;

    covergroup {
        option.at_least = 2;
        c: coverpoint color;
    } cs1_inst;
}
```

Example 201—Creating an in-line covergroup instance

15.3 Defining coverage points

A **covergroup** can contain one or more coverage points. A coverage point specifies an integer expression or **enum** that is to be covered. Each coverage point includes a set of bins associated with the sampled values of the covered expression. The bins can be explicitly defined by the user or automatically created by the PSS processing tool. The syntax for specifying coverage points is shown in Syntax 71.

Evaluation of the coverage point expression (and of its enabling **iff** condition, if any) takes place when the **covergroup** is sampled (see 15.6).

15.3.1 Syntax

The syntax for **coverpoint**s is shown in Syntax 71.

Copyright © 2026 Accellera. All rights reserved.

```
covergroup_coverpoint ::= [ [ data_type ] coverpoint_identifier : ] coverpoint
   expression [ iff ( expression ) ] bins_or_empty
bins_or_empty ::=
    { { covergroup_coverpoint_body_item } }
   | ;
covergroup_coverpoint_body_item ::=
    covergroup_option
   | covergroup_coverpoint_binspec
```

Syntax 71—Coverpoint declaration

The following also apply:

a) A **coverpoint** coverage point creates a hierarchical scope and can be optionally labeled. The label (*coverpoint_identifier*) designates the name of the coverage point. This name can be used to add this coverage point to a cross coverage specification. If the coverage point is associated with a single variable and the label is omitted, the variable name becomes the name of the coverage point. A coverage point on an expression is required to specify a label.

b) A data type for the coverpoint may be specified. The data type shall be an integer or **enum** type. If a data type is specified, then a label shall also be specified.

c) If a data type is specified, the **coverpoint** expression shall be assignment compatible with the data type. Values for the **coverpoint** shall be of the specified data type and shall be determined as though the **coverpoint** expression were assigned to a variable of the specified type.

d) If no data type is specified, the inferred type for the **coverpoint** shall be the self-determined type of the **coverpoint** expression.

e) The expression within the **iff** construct specifies an optional condition that disables coverage sampling for that **coverpoint**. If the **iff** *expression* evaluates to *false* at a sampling point, the coverage point is not sampled.

f) A coverage point bin associates a name and a count with a set of values. The count is incremented every time the coverage point matches one of the values in the set. The bins for a coverage point can be defined using the **bins** construct to name each bin. If the **bins** are not explicitly defined, they are automatically created by the PSS processing tool. The number of automatically created bins can be controlled using the **`auto_bin_max`** coverage option. Coverage options are described in Table 25.

15.3.2 Example

In Example 202, coverage point `s0` is covered only if `is_s0_enabled` is *true*.

```
struct s {
   rand bit[4] s0;
   rand bool   is_s0_enabled;

   covergroup {
      coverpoint s0 iff (is_s0_enabled);
   } cs4;
}
```

Example 202—Specifying an iff condition

Copyright © 2026 Accellera. All rights reserved.

15.3.3 Specifying bins

The **bins** construct creates a separate bin for each value in the given range list or a single bin for the entire range of values. The syntax for defining bins is shown in Syntax 72.

15.3.3.1 Syntax

The syntax for **bin**s is shown in Syntax 72.

```
covergroup_coverpoint_binspec ::= bins_keyword identifier
  [ [ [ constant_expression ] ] ] = coverpoint_bins
coverpoint_bins ::=
    [ covergroup_range_list ] [with ( covergroup_expression ) ] ;
  | coverpoint_identifier with ( covergroup_expression ) ;
  | default ;
covergroup_range_list ::= covergroup_value_range { , covergroup_value_range }
covergroup_value_range ::=
     expression
  | expression .. [ expression ]
  | [ expression ] .. expression
bins_keyword ::= bins | illegal_bins | ignore_bins
covergroup_expression ::= expression
```

Syntax 72—Bins declaration

The following also apply:

a) To create a separate bin for each value (an array of bins), add square brackets (**[ ]**) after the bin name.
   1) To create a fixed number of bins for a set of values, a single positive integral expression can be specified inside the square brackets.
   2) The bin name and optional square brackets are followed by a *covergroup_range_list* that specifies the set of values associated with the bin.
   3) It shall be legal to use the range value form *expression*.. and ..*expression* to denote a range that extends to the upper or lower value (respectively) of the coverpoint data type.

b) If a fixed number of bins is specified and that number is smaller than the specified number of values, the possible bin values are uniformly distributed among the specified bins.
   1) The first *N* specified values (where *N* = int(number of values / number of bins)) are assigned to the first bin, the next *N* specified values are assigned to the next bin, etc.
   2) Duplicate values are retained; thus, the same value can be assigned to multiple bins.
   3) If the number of values is not evenly divisible by the number of bins, then the last bin will include the remaining items, e.g., for

      ```
      bins fixed [4] = [1..10, 1, 4, 7];
      ```

      The 13 possible values are distributed as follows: `<1,2,3>`, `<4,5,6>`, `<7,8,9>`, `<10,1,4,7>`.

c) A *covergroup_expression* is an *expression*. In the case of a **with** *covergroup_expression*, the expression can involve constant terms and the **coverpoint** variable (see 15.3.3.3).

Copyright © 2026 Accellera. All rights reserved.

d) The **default** specification defines a bin that catches the values of the coverage point that do not lie within any of the defined bins. The default is useful for catching unplanned or invalid values. The coverage calculation for a coverage point shall not take into account the coverage captured by default bins. Default bins are also excluded from cross coverage (see 15.4). A default bin cannot be explicitly ignored (see 15.3.5).

15.3.3.2 Example

In Example 203, the first **bins** construct associates bin `a` with the values of `v_a`, between `0` and `63` and the value `65`. The second **bins** construct creates a set of 65 bins `b[127]`, `b[128]`, … `b[191]`. Note that when empty square brackets are specified, each value is assigned one bin, including values that are specified more than once. Likewise, the third **bins** construct creates 3 bins: `c[200]`, `c[201]`, and `c[202]`. The fourth **bins** construct associates bin `d` with the values between `1000` and `1023` (the trailing `..` represents the maximum value of `v_a`). Every value that does not match bins `a`, `b[]`, `c[]`, or `d` is added into its own distinct bin (e.g., the value `64`), using the **default** specification.

```
struct s {
   rand bit[10] v_a;

   covergroup {
      coverpoint v_a {
         bins a   = [0..63, 65];
         bins b[] = [127..150, 148..191];
         bins c[] = [200, 201, 202];
         bins d   = [1000..];
         bins others[] = default;
      }
   } cs;
}
```

Example 203—Specifying bins

15.3.3.3 Coverpoint bin with covergroup expressions

The **with** clause specifies that only those values in the *covergroup_range_list* (see Syntax 72) that satisfy the given expression (i.e., for which the expression evaluates to *true*) are included in the bin. In the expression, the name of the **coverpoint** shall be used to represent the candidate value. The candidate value is of the same type as the **coverpoint**.

The **with** clause behaves as if the expression were evaluated for every value in the *covergroup_range_list* at the time the covergroup instance is created. The **with** *covergroup_expression* is applied to the set of values in the *covergroup_range_list* prior to distribution of values to the bins. The result of applying a **with** *covergroup_expression* shall preserve multiple, equivalent bin items as well as the bin order. The intent of these rules is to allow the use of non-simulation analysis techniques to calculate the bin (e.g., formal symbolic analysis) or for caching of previously calculated results.

Consider Example 204, where the bin definition selects all values from `0` to `255` that are evenly divisible by `3`.

Copyright © 2026 Accellera. All rights reserved.

```
struct s {
    rand bit[8] x;

    covergroup {
        a: coverpoint x {
              bins mod3[] = [0..255] with ((a % 3) == 0);
        }
    } cs;
}
```

Example 204—Select constrained values between 0 and 255

The name of the **coverpoint** itself may be used in place of the *covergroup_range_list*, preceding the **with** keyword, to denote all values of the **coverpoint**. Only the name of the **coverpoint** containing the bin being defined shall be allowed.

In Example 205, **coverpoint** name `a` is used in place of the *covergroup_range_list* to denote that the **with** *covergroup_expression* will be applied to all values of the **coverpoint**.

```
struct s {
    rand bit[8] x;

    covergroup {
        a: coverpoint x {
              bins mod3[] = a with ((a % 3) == 0);
        }
    } cs;
}
```

Example 205—Using with in a coverpoint

15.3.4 Automatic bin creation for coverage points

If a coverage point does not define any bins, PSS automatically creates bins. This provides an easy-to-use mechanism for binning different values of a coverage point. Users can either let the tool automatically create bins for coverage points or explicitly define named bins for each coverage point.

When the automatic bin creation mechanism is used, PSS creates *N* bins to collect the sampled values of a coverage point. The value *N* is determined as follows:

— For an **enum** coverage point, *N* is the cardinality of the enumeration.

— For an integer coverage point, *N* is the minimum of 2^M and the value of the **`auto_bin_max`** option (see Table 25), where *M* is the number of bits needed to represent the coverage point.

If the number of automatic bins is smaller than the number of possible values ($N < 2^M$), the 2^M values are uniformly distributed in the *N* bins. If the number of values, 2^M, is not divisible by *N*, then the last bin will include the additional remaining items. For example, if *M* is `3` and *N* is `3`, the eight possible values are distributed as follows: `<0..1>`, `<2..3>`, `<4..7>`.

PSS implementations can impose a limit on the number of automatic bins. See Table 25 for the default value of **`auto_bin_max`**.

Each automatically created bin will have a name of the form **`auto[`***`value`***`]`**, where *`value`* is either a single coverage point value or the range of coverage point values included in the bin (in the form

Copyright © 2026 Accellera. All rights reserved.

low..high). For enumeration types, *value* is the named constant (enum item) associated with the particular enumeration value.

15.3.5 Excluding coverage point values

A set of values associated with a coverage point can be explicitly excluded from coverage by specifying them as **ignore_bins**. See Example 206.

All values associated with ignored bins are excluded from coverage. Each ignored value is removed from the set of values associated with any coverage bin. The removal of ignored values shall occur after distribution of values to the specified bins.

Example 206 may result in a bin that is associated with no values or sequences. Such empty bins are excluded from coverage.

```
struct s {
    rand bit[4] a;

    covergroup {
        coverpoint a {
            ignore_bins ignore_vals = [7, 8];
        }
    } cs23;
}
```

Example 206—Excluding coverage point values

15.3.6 Specifying illegal coverage point values

A set of values associated with a coverage point can be marked as illegal by specifying them as **illegal_bins**. See Example 207.

All values associated with illegal bins are excluded from coverage. Each illegal value is removed from the set of values associated with any coverage bin. The removal of illegal values shall occur after the distribution of values to the specified bins. If an illegal value occurs, a runtime error shall be issued. Illegal bins take precedence over any other bins, i.e., they result in a runtime error even if they are also included in another bin.

Example 207 may result in a bin that is associated with no values or sequences. Such empty bins are excluded from coverage.

```
struct s {
    rand bit[4] a;

    covergroup {
        coverpoint a {
            illegal_bins illegal_vals = [7, 8];
        }
    } cs23;
}
```

Example 207—Specifying illegal coverage point values

Copyright © 2026 Accellera. All rights reserved.

15.3.7 Value resolution

A *coverpoint expression*, the expressions in a **bins** construct, and the **coverpoint** type, if present, are all involved in comparison operations in order to determine into which bins a particular value falls. Let *e* be the coverpoint expression and *b* be an expression in a **bins** *covergroup_range_list*. The following rules shall apply when evaluating *e* and *b*:

a) If there is no coverpoint type, the effective type of *e* shall be self-determined. In the presence of a coverpoint type, the effective type of *e* shall be the coverpoint type.

b) *b* shall be statically cast to the effective type of *e*. An implementation shall issue a warning under the following conditions:

   1) If the effective type of *e* is unsigned and *b* is signed with a negative value.

   2) If assigning *b* to a variable of the effective type of *e* would yield a value that is not equal to *b* under normal comparison rules for ==.

If a warning is issued for a **bins** element, the following rules shall apply:

— If an element of a bins *covergroup_range_list* is a singleton value *b*, that element shall not appear in the bins values.

— If an element of a bins *covergroup_range_list* is a range `b1..b2` and there exists at least one value in the range for which a warning would not be issued, the range shall be treated as containing the intersection of the values in the range and the values expressible by the effective type of *e*.

Example 208 leads to the following:

— For `b1`, a warning is issued for the range `6..10`. `b1` is treated as though it had the specification `[1, 2..5, 6..7]`.

— For `b2`, a warning is issued for the range `1..10` and for the values `-1` and `15`. `b2` is treated as though it had the specification `[1..7]`.

— For `b3`, a warning is issued for the ranges `2..5` and `6..10`. `b3` is treated as though it had the specification `[1, 2..3]`.

— For `b4`, a warning is issued for the range `1..10` and for the value `15`. `b4` is treated as though it had the specification `[-1, 1..3]`.

```
struct s {
   rand bit[3] p1;        // type expresses values in the range  0 to 7
   int [3]     p2;        // type expresses values in the range -4 to 3

   covergroup {
      coverpoint p1 {
         bins b1 = [1, 2..5, 6..10]; // warning issued for range 6..10
         bins b2 = [-1, 1..10, 15];  // warning issued for range 1..10
      }                              //           and values -1 and 15
      coverpoint p2 {
         bins b3 = [1, 2..5, 6..10]; // warning issued for ranges 2..5
                                     //                       and 6..10
         bins b4 = [-1, 1..10, 15];  // warning issued for range  1..10
      }                              //                     and value 15
   } c1;
}
```

Example 208—Value resolution

Copyright © 2026 Accellera. All rights reserved.

15.4 Defining cross coverage

A **covergroup** can specify cross coverage between two or more coverage points or variables. Cross coverage is specified using the **cross** construct (see Syntax 73). When a variable `V` is part of a cross coverage, the PSS processing tool shall implicitly create a coverage point for the variable, as if it had been created by the statement `coverpoint V;`. Thus, a *cross* involves only coverage points. Expressions cannot be used directly in a **cross**; a coverage point shall be explicitly defined first.

15.4.1 Syntax

Syntax 73 declares a **cross**.

```
covergroup_cross ::= covercross_identifier : cross
  coverpoint_identifier { , coverpoint_identifier }
  [iff ( expression )] cross_item_or_null
cross_item_or_null ::=
    { { covergroup_cross_body_item } }
  | ;
covergroup_cross_body_item ::=
    covergroup_option
  | covergroup_cross_binspec
covergroup_cross_binspec ::=
  bins_keyword identifier = covercross_identifier with ( covergroup_expression ) ;
covergroup_expression ::= expression
```

Syntax 73—Cross declaration

The following also apply:

a) The label is required for a **cross**.

b) The expression within the optional **iff** provides a conditional sampling guard for the cross coverage. If the condition evaluates to *false* at any sampling point, the cross coverage is not sampled.

c) Cross coverage of a set of *N* coverage points is defined as the coverage of all combinations of all bins associated with the N coverage points, i.e., the Cartesian product of the *N* sets of coverage point bins. See also Example 209.

15.4.2 Example

The covergroup `cov` in Example 209 specifies the cross coverage of two 4-bit variables, `a` and `b`. The PSS processing tool implicitly creates a coverage point for each variable. Each coverage point has 16 bins, specifically `auto[0]..auto[15]`. The cross of `a` and `b` (labeled `aXb`), therefore, has 256 cross products and each cross product is a bin of `aXb`.

Copyright © 2026 Accellera. All rights reserved.

```
struct s {
    rand bit[4] a, b;

    covergroup {
        aXb : cross a, b;
    } cov;
}
```

Example 209—Specifying a cross

15.4.3 Defining cross bins

In addition to specifying the coverage points that are crossed, PSS allows the definition of cross coverage bins. Cross coverage bins are specified to group together a set of cross products. A *cross coverage bin* associates a name and a count with a set of cross products. The count of the bin is incremented any time any of the cross products match; i.e., every coverage point in the **cross** matches its corresponding bin in the cross product.

User-defined bins for cross coverage are defined using **bins with** expressions. The names of the **coverpoint**s used as elements of the cross coverage are used in the **with** expressions. User-defined cross bins and automatically generated bins can coexist in the same **cross**. Automatically generated bins are retained for those cross products that do not intersect cross products specified by any user-defined cross bin.

```
struct s {
    rand bit[8] a, b;
    covergroup {
      cp_a : coverpoint a {
        bins low[] = [0..127];
        bins med = [128..143];
        bins high = [144..255];
      }
      cp_b : coverpoint b {
        bins even[] = b with (b%2 == 0);
      }
      X : cross cp_a, cp_b {
        bins small_a_b = X with (cp_a<=10 && cp_b<=10);
        bins med_a_b = X with (cp_a==128 && cp_b<=10);
        bins med_a_b_iff = X with (cp_a==128 && cp_b<=10) iff (a==128);
      }
    } cov;
}
```

Example 210—Specifying cross bins

Consider Example 210, where two coverpoints are declared for fields `a` and `b`. A cross coverage is specified between these two coverpoints. The `small_a_b` bin collects those bins where both `a<=10` and `b<= 10`.

In Example 210, bins `med_a_b` are hit when the value of "`a`" selects the bin(s) containing `128` and the value of "`b`" selects `a` bin(s) containing `b<=10`. `a==128,b==2` is one example of such a bin. `a==129,b==2` is as well. Bin `med_a_b_iff` will only be covered when `a==128` due to the iff sampling condition. Specifically:

- bins `small_a_b` is hit when `low[0]..low[10]` and `even[0]..even[5]` are hit
- bins `med_a_b` is hit when the value of `a` is within the `med` bin and `even[0]..even[5]` are hit

Copyright © 2026 Accellera. All rights reserved.

— bins `med_a_b_iff` is hit when the value of `a` is within the `med` bin, `even[0]…even[5]` are hit, and the variable `a` is `128`. Note that the iff excludes `a==129` from causing this bin to be hit.

15.5 Specifying coverage options

Options control the behavior of the **covergroup**, **coverpoint**, and **cross** elements. Options can be specified when creating an instance of a reusable **covergroup**, and are specific to that **covergroup** instance.

Specifying a value for the same option more than once within the same **covergroup** definition shall be an error. Specifying a value for the option more than once when creating a **covergroup** instance shall be an error.

Table 25 lists the instance-specific **covergroup** options and their description. Each instance of a reusable **covergroup** type can initialize an instance-specific option to a different value.

Table 25—Instance-specific covergroup options

Option name	Default	Description
weight=number	1	If set at the **covergroup** syntactic level, it specifies the weight of this **covergroup** instance relative to all other instances when computing overall instance coverage. If set at the **coverpoint** (or **cross**) syntactic level, it specifies the weight of a **coverpoint** (or **cross**) for computing the instance coverage of the enclosing **covergroup**. The specified weight shall be a non-negative integral value.
goal=number	100	Specifies the target goal for a **covergroup** instance or for a **coverpoint** or **cross**. The specified value shall be a non-negative integral value.
name=string	unique name	Specifies a name for the **covergroup** instance. If unspecified, a unique name for each instance shall be automatically generated by the tool.
comment=string	""	A comment that appears with the **covergroup** instance or with a **coverpoint** or **cross** of a **covergroup** instance. The comment is saved in the coverage database and included in the coverage report.
at_least=number	1	Minimum number of hits for each bin. A bin with a hit count that is less than *number* is not considered covered. The specified value shall be a positive integral value.
detect_overlap=bool	false	When *true*, a warning is issued if there is an overlap between the range list of two bins of a **coverpoint**.
auto_bin_max=number	64	Maximum number of automatically created bins when no bins are explicitly defined for a **coverpoint**. The specified value shall be a positive integral value.
per_instance=bool	false	Each instance contributes to the overall coverage information for the **covergroup** type. When *true*, coverage information for this **covergroup** instance shall be saved in the coverage database and included in the coverage report. When *false*, implementations are not required to save instance-specific information.

Copyright © 2026 Accellera. All rights reserved.

Instance options can be specified at the **covergroup** level. Except for the **`weight`**, **`goal`**, **`comment`**, and **`per_instance`** options (see Table 25), all other options set at the covergroup syntactic level act as a default value for the corresponding option of all **coverpoint**s and **cross**es in the **covergroup**. Individual **coverpoint**s and **cross**es can overwrite these defaults. When set at the **covergroup** level, the **`weight`**, **`goal`**, **`comment`**, and **`per_instance`** options do not act as default values to the lower syntactic levels.

15.5.1 Example

The instance-specific options mentioned in Table 25 can be set in the **covergroup** definition. Example 211 shows this, and how coverage options can be set on a specific **coverpoint**.

```
covergroup cs1 (bit[64] a_var, bit[64] b_var) {
   option.per_instance = true;
   option.comment = "This is CS1";

   a : coverpoint a_var {
      option.auto_bin_max = 128;
   }

   b : coverpoint b_var {
      option.weight = 10;
   }
}
```

Example 211—Setting options

15.6 Covergroup sampling

Coverage credit can be taken once execution of the **action** containing **covergroup** instance(s) is complete. Thus, by default, all **covergroup** instances that are created as a result of a given **action**'s traversal are sampled when that **action**'s execution completes. **covergroup** sampling in monitors is described in 16.5.1. Table 26 summarizes when **covergroup**s are sampled, based on the context in which they are instantiated.

Table 26—Covergroup sampling

Instantiation context	Sampling point
Flow objects	Sampled when the outputting **action** completes traversal.
Resource objects	Sampled before the first **action** referencing them begins traversal.
Action	Sampled when the instantiating **action** completes traversal.
Monitor	Sampled at the match point of the cover statement, instantiating the monitor.
Data structures	Sampled along with the context in which the data structure is instantiated, e.g., if a data structure is instantiated in an **action**, the **covergroup** instantiated in the data structure is sampled when the **action** completes traversal.

15.7 Per-type and per-instance coverage collection

By default, **covergroup**s collect coverage on a *per-type* basis. This means that all coverage values sampled by instances of a given **covergroup** type, where **`per_instance`** is *false*, are merged into a single collection.

Copyright © 2026 Accellera. All rights reserved.

Per-instance coverage is collected when **`per_instance`** is *true* for a given **covergroup** instance and when a contiguous path of named handles exists from the root component, root action, or an instantiated cover statement to where new instances of the containing type are created. If one of these conditions is not satisfied, *per-type* coverage is collected for the **covergroup** instance.

15.7.1 Per-instance coverage of flow and resource objects

Per-instance coverage of flow objects (**buffer** (see 9.3.1), **stream** (see 9.3.2), **state** (see 9.3.3)) and **resource** objects (see 9.4.1)) is collected for each pool of that type.

In Example 212, there is one pool (`pss_top.b1_p`) of buffer type `b1`. When the PSS model runs, coverage from all 10 executions of `P_a` and `C_a` is placed in the same coverage collection that is associated with the pool through which `P_a` and `C_a` exchange the buffer object `b1`.

```
enum mode_e { M0, M1, M2 }

buffer b1 {
    rand mode_e mode;

    covergroup {
        option.per_instance = true;

        coverpoint mode;
    } cs;
}

component pss_top {
    pool b1 b1_p;
    bind b1_p *;

    action P_a {
        output b1 b1_out;
    }

    action C_a {
        input b1 b1_in;
    }

    action entry {
        activity {
            repeat (10) {
                do C_a;
            }
        }
    }
}
```

Example 212—Per-instance coverage of flow objects

15.7.2 Per-instance coverage in actions

Per-instance coverage for **action**s is enabled when **`per_instance`** is *true* for a **covergroup** instance and when a contiguous path of named handles exists from the root action to the location where the **covergroup** is instantiated.

Copyright © 2026 Accellera. All rights reserved.

In Example 213, a contiguous path of named handles exists from the root action to the covergroup instance inside `a1` (`entry.a1.cg`). Coverage data collected during traversals of action `A` are placed in a coverage collection unique to this named path. Plus, four samples are placed in the coverage collection associated with the instance path `entry.a1.cg` because the named action handle `a1` is traversed four times.

Also in Example 213, a contiguous path of named handles does not exist from the root action to the covergroup instance inside the action traversal by type (`do A`). In this case, coverage data collected during the 10 traversals of action `A` by type (`do A`) are placed in the per-type coverage collection associated with covergroup type `A::cg`.

```
enum mode_e { M0, M1, M2 }

component pss_top {

    action A {
        rand mode_e mode;

        covergroup {
            option.per_instance = true;

            coverpoint mode;
        } cg;
    }

    action entry {
        A      a1;
        activity {
            repeat (4) {
                a1;
            }
            repeat (10) {
                do A;
            }
        }
    }
}
```

Example 213—Per-instance coverage in actions

Copyright © 2026 Accellera. All rights reserved.

16. Behavioral coverage

A large number of specific scenarios may be generated from one PSS specification. These scenarios vary in the action order and data. Coverage statements identify a key action order and data combinations that need to be observed to exercise the key functionality. The specification of the observed action order and, possibly their data, related to this ordering, is called *behavioral coverage* and is described in this clause. Data coverage is described in Clause 15.

16.1 Defining behavioral coverage: cover and monitor

The **cover** statement directs the PSS processing tool to observe an action scenario; i.e., an action execution order together with its data as specified in its body. The action order specification, including action data specification, is called a scenario. Scenarios may also be written separately as monitors. The cover statements are only active in component instances actually instantiated from the root component.

The syntax for monitors and cover statements is shown in Syntax 74.

```
cover_stmt ::=
  [ label_identifier : ] cover type_identifier ;
| [ label_identifier : ] cover { { monitor_body_item } }
monitor_declaration ::= monitor monitor_identifier
  [ template_param_decl_list ] [ monitor_super_spec ] { { monitor_body_item  } }
abstract_monitor_declaration ::= abstract monitor_declaration
monitor_super_spec ::= : type_identifier
monitor_body_item ::=
    monitor_activity_declaration
  | override_declaration
  | monitor_constraint_declaration
  | monitor_field_declaration
  | covergroup_declaration
  | attr_group
  | compile_assert_stmt
  | covergroup_instantiation
  | monitor_body_compile_if
  | annotation_usage monitor_body_item
  | stmt_terminator
monitor_field_declaration ::=
    const_field_declaration
  | action_handle_declaration
  | monitor_handle_declaration
```

Syntax 74—Cover statement and monitor declaration

Monitors are coverage counterparts of actions. An action construct describes a generation scenario for a solution. A monitor construct describes a scenario to be observed, and a cover statement instructs the PSS processing tool to monitor the executed stream of actions for the presence of the specified scenario.

Copyright © 2026 Accellera. All rights reserved.

A **cover** statement may incorporate a monitor type directly or instantiate an existing **monitor** type, as shown in Example 214.

A **monitor** type may be declared separately using the **monitor** keyword and a *monitor_identifier*, as shown in Syntax 74. The monitor's syntax is similar, but not identical to the syntax of the compound action (see 10 and 11). Like an action, a monitor may declare fields, have an activity, constraints, and covergroup declarations and instantiations. All rules applicable to actions apply also to monitors unless the opposite is stated explicitly.

A **monitor** declaration optionally specifies a *monitor_super_spec*, a previously defined monitor type from which the new type inherits its members. A **monitor** activity specifies the scenario that shall be observed in order to satisfy the monitor. If more than one activity is specified in a monitor, its scenario is equivalent to its activity scenarios combined in a schedule of scenarios (see 16.3.8).

The following also apply:

a) Monitor fields may be action and monitor handles. Data attributes, references, and resource claims are not supported in monitors. Other data fields shall be declared as **static const**.

b) An *abstract monitor* may be declared as a template that defines a base set of field attributes from which other monitors may inherit. Non-abstract derived monitors may be instantiated like any other monitor. Abstract monitors shall not be instantiated directly.

c) An abstract monitor may be derived from another abstract monitor but not from a non-abstract monitor.

d) Abstract monitors may be extended, but the monitor remains abstract and may not be instantiated directly.

A non-abstract monitor (the initial definition and all its extensions) shall have one or more **activity** statements (see 16.3).

An abstract monitor may have no activity defined.

In Example 214, the **cover** statement `c1` captures a scenario when the execution of an action of type `read` follows the execution of an action of type `write`. Cover statement `c1` directly specifies the monitor scenario as part of the cover directive. Cover statement `c2` specifies the scenario as a reusable monitor type `wr`.

Copyright © 2026 Accellera. All rights reserved.

```
component pss_top {

  action read {}
  action write {}

  c1: cover {
    activity {
      do write;
      do read;
    }
  }

  monitor wr {
    activity {
      do write;
      do read;
    }
  }

  c2: cover wr; // equivalent to c1
}
```

Example 214—Cover statement and monitor

16.2 Behavioral coverage concepts

This section defines basic concepts used in the definition of the behavioral coverage.

a) A *top-level monitor* is the **monitor** type (inlined or instantiated) of a **cover** statement.

b) An observed *scenario* is a mapping of a monitor or of a monitor **activity** statement into sets of observed action executions.

c) The time is assumed to be an unsigned integer. An action execution spans from its `start` time (including) and its end time (excluding). Hence, an instantaneous action with the `start` time *t* has the end time *t+1*.

d) A *scenario checkpoint* is a time instant relative to which the scenario is observed.

e) An *attempt* is a scenario along with its checkpoint.

f) An *attempt scenario realization* is a set of action executions traversed by this attempt along with the mapping of action handles into specific action executions.

g) Realizations of an attempt of a scenario with a standalone constraint are those realizations of the unconstrained scenario that satisfy the constraint.

h) Realizations of an attempt of a scenario with a constraint, whether standalone or in-line, shall satisfy all constraints.

i) A *scenario realization start* (or beginning) *point* is a `start` time of the first observed action execution of this scenario for a given checkpoint. It may either coincide with the checkpoint, or it may be after it.

j) A *scenario realization endpoint* or *match point* is the end time of the last observed action execution of this scenario for a given checkpoint.

k) An attempt of the top-level **monitor** *has a match* or *is successful* if it has at least one scenario realization whose `start` point coincides with the attempt's checkpoint; otherwise, the attempt *has no match* or is *failing*.

l) The *first match* of a successful top-level **monitor** attempt is the closest among the match points of its scenario realizations to its checkpoint (not to the realization's `start` point).

Copyright © 2026 Accellera. All rights reserved.

m) The top-level **monitor** of a **cover** statement shall not have attempts with an empty realization (see 16.3.7).

n) A **cover** statement *has a match* or is *successful* if its top-level **monitor** has at least one successful attempt.

The notion of a checkpoint is required to build compound monitors from the smaller ones. For example, if the top-level **monitor** is a sequence of two action traversals, then the first action traversal is matched from its `start` point; whereas, the second action traversal should be matched not from its `start` point but from the end point of the first action traversal, i.e., the checkpoint of the second action traversal is the endpoint of the first action traversal. This is explained in Example 215 and the diagrams shown in Figure 21 and Figure 22.

```
action read {}
action write {}
action idle {}
action send {}
action receive {}

monitor m1 {
  write w;
  read r;
  activity {
    w;
    r;
  }
}

monitor m2 {
  activity {
    do write;
    select {
      do read;
      do send;
    };
    do receive;
  }
}

monitor m3 {
  activity {
    select {
      do write;
      do read;
    };
    select {
      do send;
      do receive;
    };
  }
}

c1: cover m1;
c2: cover m3;
```

Example 215—Illustration of behavioral coverage concepts

Copyright © 2026 Accellera. All rights reserved.

The monitors `m1`, `m2`, and `m3` define the following scenarios:

- Monitor `m1`: a `write` action execution is followed (directly or not) by a `read` action execution.
- Monitor `m2`: a `write` action execution is followed (directly or not) by either a `read` or a `send` action execution (or both), followed by a `receive` action execution.
- Monitor `m3`: An execution of a `write` or a `read` action (or both) is followed by an execution of a `send` or a `receive` action (or both).

The top-level **monitor** of the **cover** statement `c1` is `m1`. Monitor `m1` describes a compound scenario. Its scenario is a sequence of two action executions defined by handles `w` and `r`. The monitor starts matching from the `start` times of all action executions, i.e., for the action trace shown in Figure 21, its checkpoint times are t_1, t_3, t_5, t_6, t_8, and t_{10}. Only attempts with `start` points t_1 and t_6 are successful, and because there exist successful attempts, the **cover** statement `c1` is successful. The `start` times of these attempts coincide with their checkpoints. The first successful attempt has two match points: t_7 (its first match) and t_{12} because both `read` actions follow `write`. There is no need to follow the top-level attempt beyond the first top match. The second successful attempt has only one match point t_{12}, which is also its first match. The first attempt has two scenario realizations: {write$_1$, read$_1$} with the mapping: `w`→`write`$_1$, `r`→`read`$_1$, and with the mapping: `w`→`write`$_1$, `r`→`read`$_2$. The second attempt has one scenario realization: {write$_2$, read$_2$}.

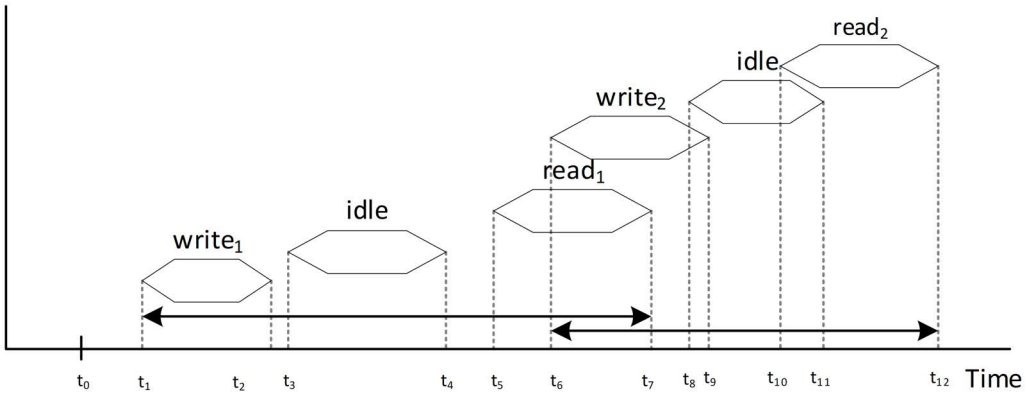


Figure 21—Monitor matching

Consider monitor `m2`, checkpoint t_0, and the trace shown in Figure 22.

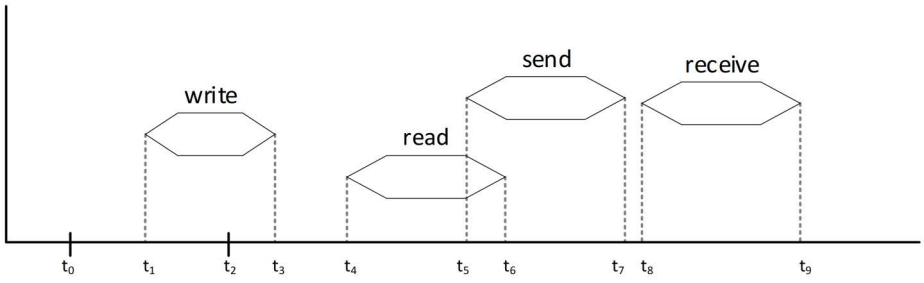


Figure 22—Behavioral coverage concepts illustration

Copyright © 2026 Accellera. All rights reserved.

This attempt has two different realizations: the action set {write, read, receive}, and the action set {write, send, receive}. Both of them have the same `start` point t_1 and the same match point t_9. Though this attempt has a single match point, it has two different realizations: the action set {write, read, receive} and the action set {write, send, receive}. The same is true for the attempt with the checkpoint t_1. For the checkpoint t_2, the attempt has no scenario realizations.

Now consider **monitor** `m3`, **cover** statement `c2`, and the trace shown in Figure 22. To check `c2`, **monitor** `m3` is matched from the beginning of every action execution, i.e., at checkpoints t_1, t_4, t_5, and t_8. There are two successful attempts, starting at times t_1 and t_4. The first attempt has scenario realizations {write, send} and {write, receive} with the `start` point t_1, and match points t_7 and t_9, correspondingly. The second attempt has one scenario realization {read, receive}, with the `start` point t_4 and the match point t_9.

16.3 Monitor activity

A scenario to be watched is defined in a **monitor** using an **activity** statement. A **monitor activity** statement is similar but not identical to an **action activity** statement.

The monitor scenario is defined by the monitor activity hierarchically by activity statements corresponding to the subscenarios.

```
monitor_activity_declaration ::=
  activity { { monitor_activity_stmt } }
monitor_activity_stmt ::=
    [ label_identifier : ] labeled_monitor_activity_stmt
  | monitor_activity_action_traversal_stmt
  | monitor_activity_monitor_traversal_stmt
  | action_handle_declaration
  | monitor_handle_declaration
  | monitor_activity_constraint_stmt
  | annotation_usage monitor_activity_stmt
  | stmt_terminator
labeled_monitor_activity_stmt ::=
    monitor_activity_sequence_block_stmt
  | monitor_activity_concat_stmt
  | monitor_activity_eventually_stmt
  | monitor_activity_overlap_stmt
  | monitor_activity_schedule_stmt
  | monitor_activity_select_stmt
  | activity_super_stmt
monitor_handle_declaration ::= monitor_type_identifier
  monitor_instantiation ;
monitor_instantiation ::=
  monitor_identifier { array_dim }
  { , monitor_identifier { array_dim } }
```

Syntax 75—Monitor activity

Copyright © 2026 Accellera. All rights reserved.

There are the following monitor activity statements for scenario specification:

- **Action** traversal scenario (see 16.3.1)
- **Sequence** statement (see 16.3.2)
- **Concat** statement (see 16.3.3)
- **Eventually** statement (see 16.3.4)
- **Overlap** statement (see 16.3.5)
- **Select** statement (see 16.3.6)
- Empty scenario (see 16.3.7)
- **Schedule** statement (see 16.3.8)
- **Monitor** traversal statement (see 16.3.9)

16.3.1 Action traversal scenario

An *action traversal scenario* specified by a monitor action traversal statement observes an execution of an action either atomic or compound. It has the same syntax as an action traversal statement in actions (see Syntax 76 and 11.3.1) with the only difference being that in this context, inline constraints shall be specified using the *monitor_constraint_set*.

```
monitor_activity_action_traversal_stmt ::= activity_action_traversal_stmt
activity_action_traversal_stmt ::=
    action_handle_traversal_stmt
  | action_type_traversal_stmt
action_handle_traversal_stmt ::=
   identifier [ [ expression ] ] [ action_initializer_list ]
  inline_constraints_or_empty
action_type_traversal_stmt ::=
   [ label_identifier : ] do type_identifier [ action_initializer_list ]
    inline_constraints_or_empty
inline_constraints_or_empty ::=
       action_activity_inline_constraints_or_empty
     | monitor_activity_inline_constraints_or_empty
monitor_activity_inline_constraints_or_empty ::=
    with monitor_constraint_set
  | ;
```

Syntax 76—Action traversal statement

identifier names a unique action handle or variable in the context of the containing monitor type or activity scope. The syntactical rules are the same as for an action traversal statement in an action activity (see 11.3.1.1).

The following also apply:

a) The semantics of traversing individual action handle array elements are the same as those of traversing individually declared action handles.

b) When traversing an action handle array as a whole, an inline constraint shall not be specified.

Copyright © 2026 Accellera. All rights reserved.

c) The anonymous action traversal statement is semantically equivalent to an action traversal with the exception that it does not create an action handle that may be referenced from elsewhere.

d) A named action handle may only be traversed once in the following scopes and nested scopes thereof:

   1) sequential activity scope (**sequence** or **concat**)

   2) overlap

   3) schedule

e) Values of action attributes mentioned in data constraints are sampled at the end of the action execution.

Given a checkpoint t_0, an action traversal scenario realization consists of an execution of an appropriately constrained action of the specified type starting at time t_1, $t_1 \geq t_0$, such that there is no other appropriately constrained action execution of this type starting at an earlier time $t: t_0 \leq t < t_1$. An action traversal scenario may have multiple realizations, as explained below (see Example 216 and Figure 26).

Example 216 defines two monitors `m1` and `m2` using anonymous action traversals and their equivalent counterparts `m11` and `m21` using action handles.

```
enum locked_e {
  LOCKED,
  UNLOCKED
};

action read {
  rand locked_e lock_mode;
}

monitor m1  {
  activity {
    do read;
  }
}

monitor m11 {
  read r;
  activity {
    r;
  }
}

monitor m2  {
  do read with lock_mode == LOCKED;
}

monitor m21 {
  read r;
  activity {
    r with lock_mode == LOCKED;
  }
}

c1: cover  m1;
```

Example 216—Action traversal in monitors

Copyright © 2026 Accellera. All rights reserved.

The monitor `m1` watches an execution of an action of type `read`. The monitor `m2` watches an execution of an action of type `read` whose lock mode is `LOCKED`.

In the action trace shown in Figure 21, **cover** statement `c1` has successful attempts starting at times t_5 and t_{11} corresponding to `read` action executions.

Now consider matching the above monitors relative to a specific checkpoint. In the below examples, the checkpoint is t_1.

First, consider matching monitor `m1` in the action trace shown in Figure 23. There is no match because there is no execution of a `read` action either at time t_1 or later.

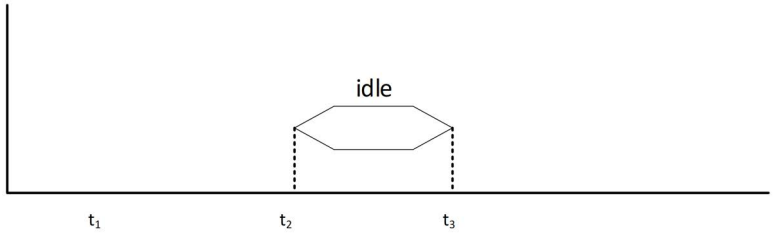


Figure 23—Action traversal statement. No match

The attempt of monitor `m1` with the checkpoint t_1 has a single scenario realization: {read} on the trace shown in Figure 24. The scenario realization `start` point is t_1 and its endpoint is t_2.

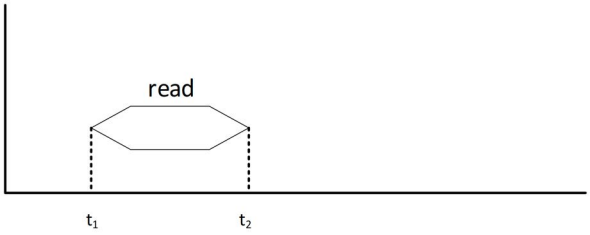


Figure 24—Action traversal statement. One scenario realization

In the trace shown in Figure 25, the monitor `m1` has one realization scenario (t_1 is a checkpoint): the unlocked `read` (the first one). The monitor `m2` also has one realization scenario: the locked `read` (the second one).

Copyright © 2026 Accellera. All rights reserved.

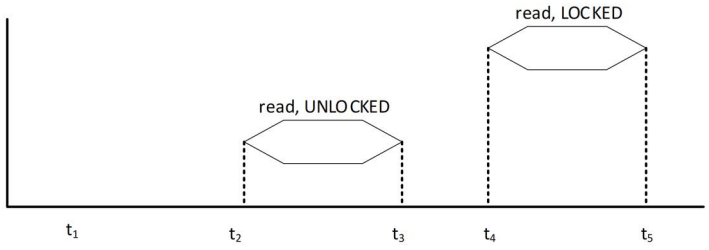


Figure 25—Matching action traversal statements with constraints

Figure 26 shows a trace where the attempt (with checkpoint t_1) of the monitor `m1` has two realization scenarios: {read$_1$} and {read$_2$}.

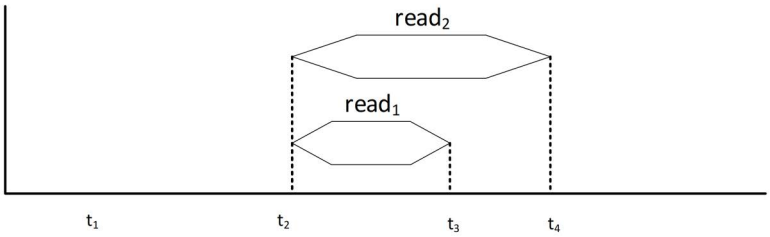


Figure 26—Action traversal, multiple matches

16.3.2 Sequential scenario

The *sequential* scenario is specified by the **sequence** statement (see Syntax 77):

```
monitor_activity_sequence_block_stmt ::= [ sequence ] {{ monitor_activity_stmt }}
```

Syntax 77—Sequence monitor activity statement

The sequential scenario defines a consecutive matching of its member-subscenarios; there may be an arbitrary gap between two consecutive subscenarios. Namely, in a sequential scenario, the first sub-scenario is matched first; at its match point or after it the second sub-scenario is matched; and so on. The realization of scenario **sequence** { $s_1, \ldots, s_n$ } consists of member-wise realization unions of subscenarios s_1, . . . , and s_n. For example, if a, b, c, d, e, and f are action executions, and for some checkpoint, a realization of s_1 is set {a,b} and a realization of s_2 for a checkpoint at or after the end of b is set {c} and for a checkpoint at or after the end of c, a realization of s_3 is set {d,e,f}, then {a,b,c,d,e,f} is a realization of sequence {s_1, s_2, s_3}.

Consider monitor `m1`:

```
monitor m1 {
  activity {
    do write;
    do read;
  }
```

Copyright © 2026 Accellera. All rights reserved.

```
}
```

defined in Example 215, the action trace shown in Figure 27, and the checkpoint t_0.

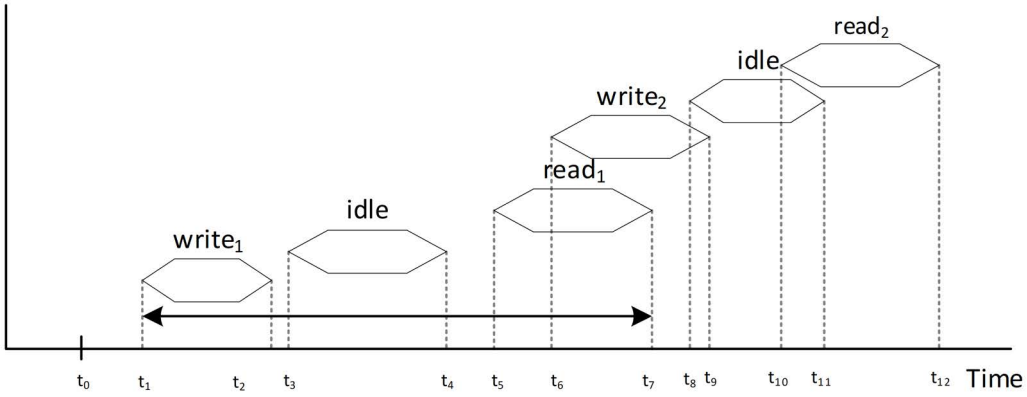


Figure 27—Sequential scenario

Monitor `m1` defines a sequence of two action traversals: {`do write; do read;`}. Its subscenarios are `do write` and `do read`. First, the scenario `do write` is matched. Its realization is the $\texttt{write}_1$ action execution, and its end time is t_2. As a checkpoint of the second scenario `do read`, t_2 or a later time instant should be chosen. For the checkpoint t_2, the second scenario realization is {read_1} action execution. For any checkpoint between t_2 and t_5 (the start point of {read_1}), the scenario realization will not change. For any checkpoint t, $t_5 < t \leq t_{10}$, the scenario realization is {read_2} action execution. For any checkpoint $t > t_{10}$, there is no match. To summarize, scenario `{ do write; do read; }` has two scenario realizations: {write_1, read_1} and {write_1, read_2}, with the match points t_7 and t_{12}, correspondingly. Informally speaking, we choose the first execution of a `write` action at or after t_0, and at or after its endpoint (t_7), we choose any `read` action.

16.3.3 Concatenation scenario

The *concatenation* scenario is specified by the monitor activity concat statement (see Syntax 78):

```
monitor_activity_concat_stmt ::= concat {{ monitor_activity_stmt }}
```

Syntax 78—Concat monitor activity statement

The concatenation scenario defines an immediate consecutive matching of its subscenarios; the checkpoint of the next sub-scenario is the matching point of the previous one. Namely, in a concatenation scenario, the first sub-scenario is matched first; at its match point, the second sub-scenario is matched, and so on. The realization of scenario `concat { `$s_1, \ldots, s_n$` }` consists of member-wise realization unions of subscenarios s_1, ..., and s_n. For example, if for some checkpoint, {a,b} is a realization of s_1, for a checkpoint at the end of `b`, {c} is a realization of s_2, and for a checkpoint at the end of `c`, {d,e,f} is a realization of s_3, then {a,b,c,d,e,f} is a realization of `concat { `s_1, s_2, s_3` }`. Here, a, b, c, d, e, and f are action executions.

The concatenation scenario often may be used interchangeably with the sequential scenario, but there are cases when the matching results are different in both scenarios (see the examples below). Because concat is

Copyright © 2026 Accellera. All rights reserved.

more restrictive, all realizations of the concatenation scenario are also realizations of the sequential scenario, but the opposite is not always true.

The following action definitions will be used in this section's examples:

```
action read {
  rand int core;
}
action write {
  rand int core;
}
action start {}
```

Example 217 shows a simple case when the sequential and concatenation scenarios behave similarly.

```
c1: cover {
  activity {
    concat {
      do write;
      do read;
    }
  }
}

c2: cover {
  activity {
    sequence {
      do write;
      do read;
    }
  }
}
```

Example 217—Concat vs sequence scenarios. Simple case

For the action trace shown in Figure 28, the results of cover statements `c1` and `c2` are identical. Both top-level scenarios have the only realization {write, read}.

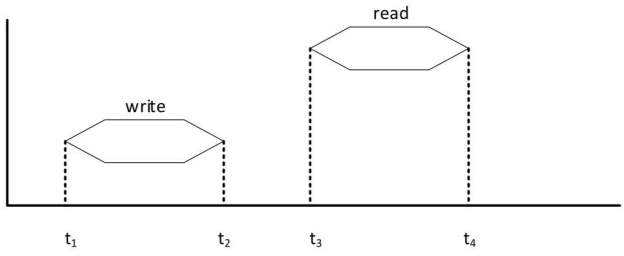


Figure 28—Action trace with two consecutive actions

In the action trace shown in Figure 29, the attempt of the **cover** statement `c1` with the checkpoint t_1 has a realization {write, $read_1$} because the match point of the first sub-scenario "`do write`" is t_2, and this time instant serves the checkpoint of the second sub-scenario `do read`. The attempt of the **cover** statement `c2` has two realizations: {write, $read_1$} and {write, $read_2$}. However, because all attempts of both of these **cover**

Copyright © 2026 Accellera. All rights reserved.

statements are either simultaneously successful or simultaneously failing, there is no difference between **sequence** and **concat** in this case as well.

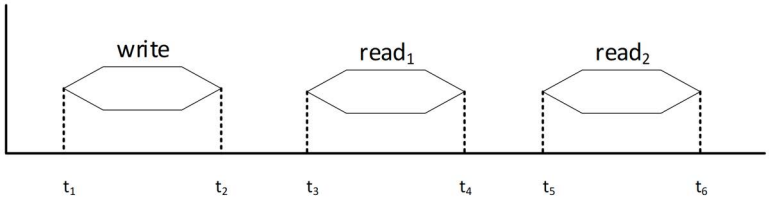


Figure 29—Action trace with two repeated actions at the end

Example 218 illustrates the minor difference between the concatenation and sequential scenarios in the presence of covergroups.

```
c3: cover {
  write w;
  activity {
    concat {
      do start;
      w;
      do read;
    }
  }
  covergroup {
    cp: coverpoint w.core;
  } cg;
}

c4: cover {
  write w;
  activity {
    sequence {
      do start;
      w;
      do read;
    }
  }
  covergroup {
    cp: coverpoint w.core;
  } cg;
}
```

Example 218—Concat vs sequence scenarios. Covergroups

In the action trace shown in Figure 30, the top-level scenario of the **cover** statement `c3` has a realization {start, write$_1$, read}; the top-level scenario of cover statement `c4` has two realizations: {start, write$_1$, read} and {start, write$_2$, read}. This difference may impact the embedded covergroup sampling: `c3` will sample `core` value 1, whereas `c4` may sample either `core` value 0 or 1 (see 16.5).

Copyright © 2026 Accellera. All rights reserved.

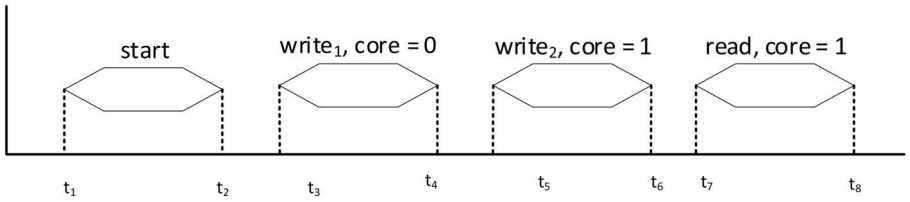


Figure 30—Action trace with two repeated actions in the middle

Example 219 and Example 220 illustrate the impact of inline and standalone constraints on the matching behavior of sequential and concatenation monitors.

In Example 219, there are two cover statements, `c5` and `c6`, each of which specifies an inline constraint for the traversal of the `write` action. The activity in `c5` specifies a concat scenario and `c6` specifies a sequence scenario.

```
c5: cover {
  activity {
    concat {
      do start;
      do write with core == 0;
      do read;
    }
  }
}

c6: cover {
  activity {
    sequence {
      do start;
      do write with core == 0;
      do read;
    }
  }
}
```

Example 219—Concat vs sequence scenarios. Inline constraints with same behavior

Consider the traces shown in Figure 31 and Figure 32, in the context of Example 219.

The top-level scenarios of `c5` and `c6` have the same realization for each trace: {start, write$_1$, read} for the trace in Figure 31 and {start, write$_2$, read} for the trace in Figure 32.

The inline constraint in the concat scenario in cover statement `c5` will match the first traversal of `write` that satisfies the constraint. In Figure 32, starting at t_2, the match point of the `start` traversal, the sub-scenario "`do write with core == 0`" has the realization {write$_2$}.

Copyright © 2026 Accellera. All rights reserved.

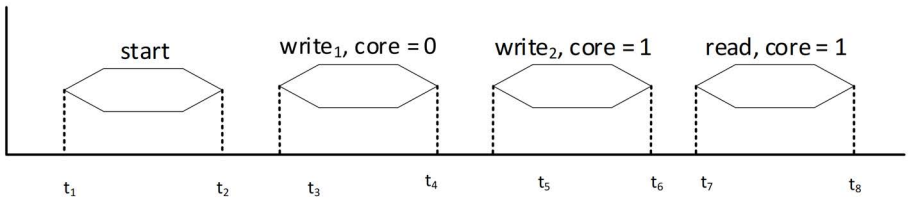


Figure 31—First alternative. Example 219

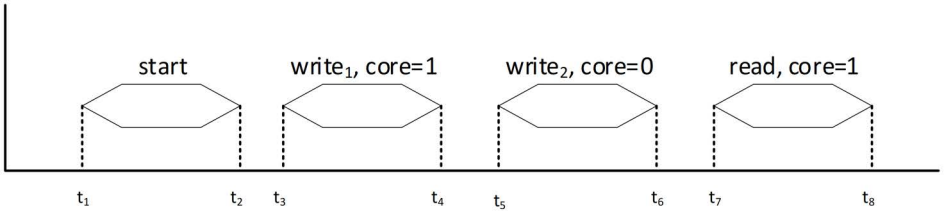


Figure 32—Second alternative. Example 219

Example 220 illustrates the case of a different behavior of the concatenation and sequential monitors.

```
c7: cover {
  write w;
  activity {
    concat {
      do start;
      w;
      do read;
    }
  }
  constraint w.core == 0;
}

c8: cover {
  write w;
  activity {
    sequence {
      do start;
      w;
      do read;
    }
  }
  constraint  w.core == 0;
}
```

Example 220—Concat vs sequence scenarios. Different behavior due to standalone constraints

In the trace shown in Figure 33, both `c7` and `c8` have the same realization {start, write$_1$, read}. In the absence of the standalone constraint, the only realization of the top-level scenario of `c7` is {start, write$_1$, read}. According to 16.2.g, the standalone constraint "`w.core == 0`" shall hold for this realization, so the realization matches. The top-level scenario of `c8` in the absence of the constraint has two

Copyright © 2026 Accellera. All rights reserved.

realizations: {start, write$_1$, read} and {start, write$_2$, read}; however, only the first realization meets the constraint condition, so that is the realization that will match.

In the trace shown in Figure 34, the top-level scenario of `c7` does not have a match since the only unconstrained realization {start, write$_1$, read} does not satisfy the constraint. The top-level scenario of `c8` does have a realization {start, write$_2$, read}.

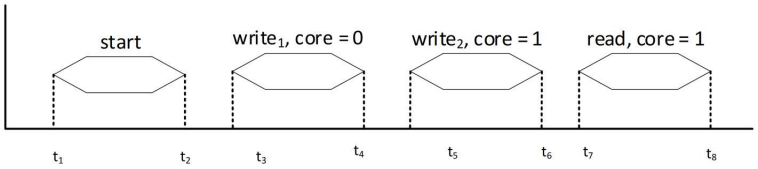


Figure 33—First alternative. Example 220

Copyright © 2026 Accellera. All rights reserved.

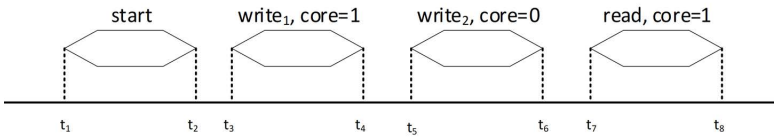


Figure 34—Second alternative. Example 220

Example 221 illustrates another difference between the concatenation and sequential scenarios in the presence of inlined constraints.

```
c9: cover {
  read r;
  write w;
  activity {
    concat {
      do start;
      w;
      r with core == w.core;
    }
  }
}

c10: cover {
  read r;
  write w;
  activity {
    sequence {
      do start;
      w;
      r with core == w.core;
    }
  }
}
```

Example 221—Concat vs sequence scenarios. Inline constraints with different behavior

The top-level scenarios of both `c9` and `c10` have the same realization {start, write$_1$, read} in the trace in Figure 35. The top-level scenario of `c10` has a realization {start, write$_2$, read} on the trace in Figure 36. The top-level scenario of `c9` does not have any realization there. Indeed, the `write` action is searched for from the checkpoint t_2. The only match is the first `write` action, and this action has attribute "`core == 0`". The `read` action is searched from the checkpoint t_4; the only `read` action on the trace has attribute "`core == 0`", and it does not satisfy the constraint "`r with core == w.core`".

Copyright © 2026 Accellera. All rights reserved.

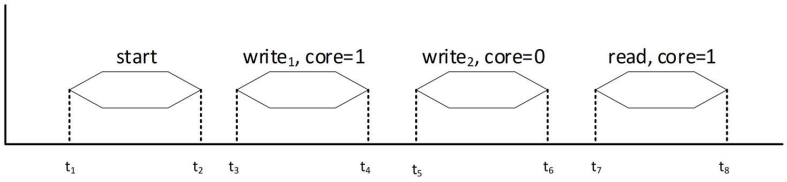


Figure 35—First alternative. Example 221

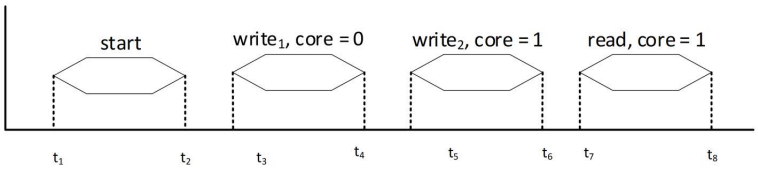


Figure 36—Second alternative. Example 221

16.3.4 Eventuality scenario

The eventuality scenario is specified by the monitor activity eventually statement (see Syntax 79):

```
monitor_activity_eventually_stmt ::= eventually monitor_activity_stmt ;
```

Syntax 79—Eventually monitor activity statement

Given a checkpoint t_0, the eventuality scenario matches its sub-scenario for any checkpoint $t \geq t_0$, i.e., any realization of the sub-scenario matched from any checkpoint $t \geq t_0$ is a realization of the eventuality scenario.

Consider **cover** statement `c` in Example 222. It reads: capture a scenario when *some* `read` action after `start` belongs to core 0 or the *first* `write` action after `start` belongs to core 1. Without eventually, only the first `read` after `start` would be checked, and if that action did not satisfy "`r.core == 0`", there is no match.

Copyright © 2026 Accellera. All rights reserved.

```
action start;
action read { rand int core; }
action write { rand int core; }
c: cover {
  read r;
  write w;
  activity {
    concat {
      do start;
      select {
        eventually r;
        w;
      }
    }
    constraint {
      w.core == 1;
      r.core == 0;
    }
  }
}
```

Example 222—Eventuality scenario

On the trace shown in Figure 37, the first alternative takes place with the realization {start, $read_2$}. On the trace shown in Figure 38, the second alternative takes place with the realization {start, write}.

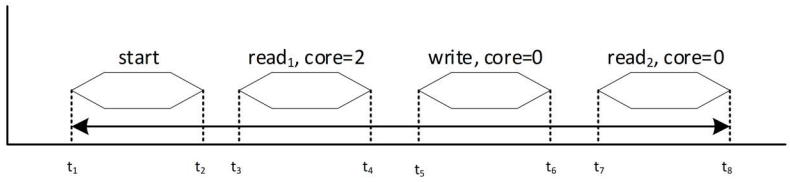


Figure 37—First alternative. Example 222

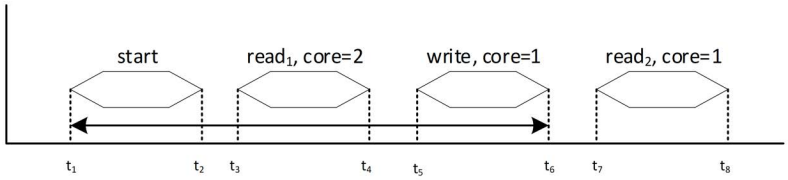


Figure 38—Second alternative. Example 222

In general, `sequence { `s_1`; `s_2`; ...; `s_n`; }` is equivalent to `concat { `s_1`; eventually `s_2`; ... eventually `s_n`; }`, where s_1, s_2, . . ., and s_n are the sequence subscenarios.

Copyright © 2026 Accellera. All rights reserved.

16.3.5 Overlapping scenario

The overlapping scenario is specified by the monitor activity **overlap** statement (see Syntax 80):

```
monitor_activity_overlap_stmt::= overlap {{ monitor_activity_stmt }}
```

Syntax 80—Overlap monitor activity statement

The overlapping scenario defines overlapping of its member subscenarios. The overlapping scenario meets the same conditions as the scheduling scenario and the additional condition that there is a time instant where all its member scenarios are simultaneously active.

In an overlapping scenario, the subscenarios are matched from the overlapping scenario checkpoint or after it. The realizations of scenario `overlap { `$s_1, \ldots, s_n$` }` consist of member-wise realization unions of subscenarios s_1, ..., and s_n, provided that these realizations of different scenarios are pairwise disjoint and that they mutually overlap in time. Realizations of scenarios s_1, ..., and s_n overlap if $\max(b_1, \ldots, b_n) < \min(e_1, \ldots, e_n)$, where b_1, ..., b_n and e_1, ..., e_n are the beginning and the end time of the realizations of scenarios s_1, ..., s_n, accordingly. For example, if {a,b} is a realization of s_1, {c} is a realization of s_2 and {d,e,f} is a realization of s_3, and the maximal among the beginning times of action executions a, c, and d is less than the minimal between the end times of action executions b, c, and f, then {a,b,c,d,e,f} is a realization of overlap `{ `s_1, s_2, s_3` }`. Here, a, b, c, d, e, and f are action executions.

Consider monitor `m` defined in Example 223, the action traces shown in Figure 39, and the checkpoint t_0.

```
action read {}
action write {}

monitor m {
  activity {
    overlap {
      do read;
      do write;
    }
  }
}
```

Example 223—Overlapping scenario

Copyright © 2026 Accellera. All rights reserved.

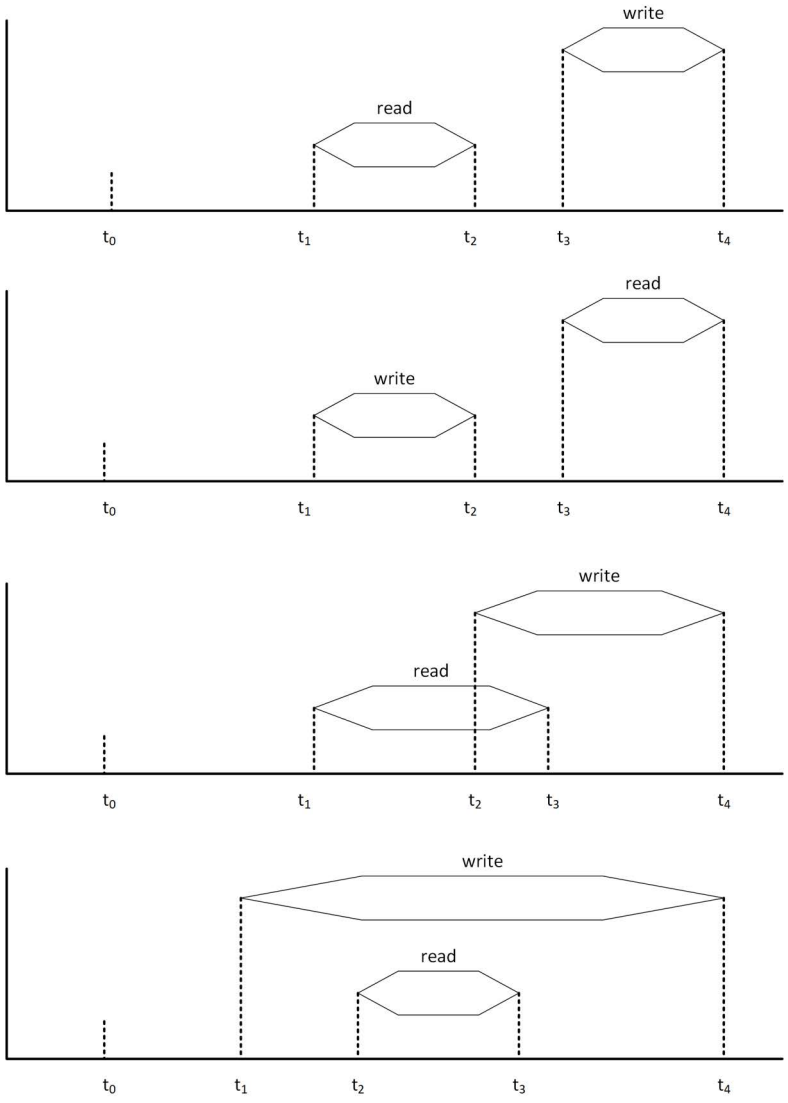


Figure 39—Overlapping scenario

The monitor's scenario has a match in the last two traces but not in the first two.

Now consider the overlapping scenario with three member scenarios defined in Example 224, the action traces shown in Figure 40 and Figure 41, and the checkpoint t_0.

Copyright © 2026 Accellera. All rights reserved.

```
action read {}
action write {}
action send {}

monitor m {
  activity {
    overlap {
      do read;
      do write;
      do send;
    }
  }
}
```

Example 224—Overlapping of three scenarios

On the trace shown in Figure 40, the overlapping scenario has one realization: {write, read, send} because the three action executions are simultaneously active (on the time interval from t_3 to t_4). On the other hand, this scenario does not have any realization on the trace shown in Figure 41 because the three action executions do not overlap, though "write" and "read" overlap, and "read" and "send" also overlap.

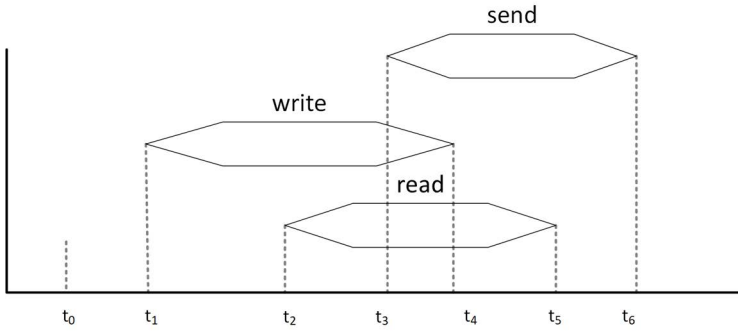


Figure 40—Overlapping of three scenarios

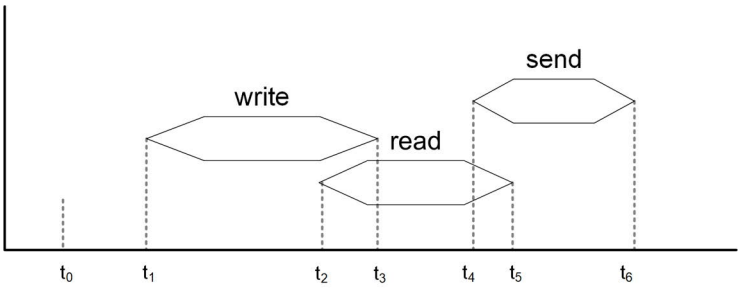


Figure 41—Three scenarios. No overlap

16.3.6 Selection scenario

The *selection* scenario is specified by the monitor activity select statement (see Syntax 81):

Copyright © 2026 Accellera. All rights reserved.

```
monitor_activity_select_stmt ::= select { monitor_activity_stmt
  monitor_activity_stmt { monitor_activity_stmt }}
```

Syntax 81—Select monitor activity statement

The selection scenario defines a selection between several alternative subscenarios. The set of realizations of the selection scenario consists of all realizations of its subscenarios. Consider Example 225 and the trace shown in Figure 42.

```
action read {}
action write {}
action idle {}
action send {}
action receive {}

c1: cover {
  activity {
    do write;
    select {
      do read;
      do send;
    };
    do receive;
  }
}

c2: cover {
  activity {
    select {
      do write;
      do read;
    };
    select {
      do send;
      do receive;
    };
  }
}
```

Example 225—Illustration of behavioral coverage concepts

Cover statement `c1` has a successful attempt starting at t_1 with two realizations: {write, read, receive} and {write, send, receive}. Cover statement `c2` has two successful attempts, one starting at time t_1 with two realizations {write, send} and {write, receive}, and the other has starting at time t_3 with the realization {read, receive}. See also Example 215.

Copyright © 2026 Accellera. All rights reserved.

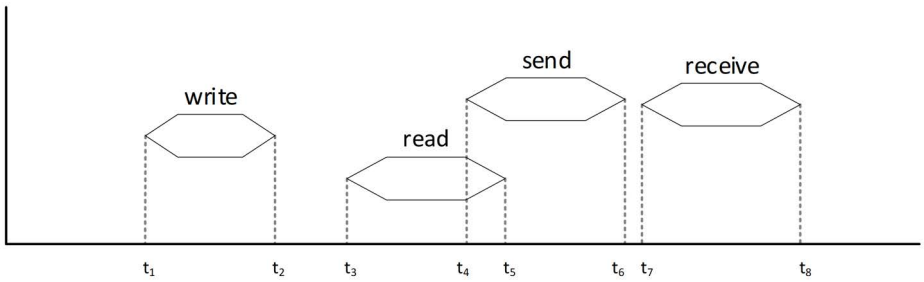


Figure 42—Illustration to Example 225

16.3.7 Empty scenario

An *empty scenario* is defined by an empty monitor, an empty activity, or an activity statement with an empty body. For example, the scenarios in Example 226 are empty:

```
monitor m {}
activity {}   // activity statement with an empty body
{}            // empty sequence
sequence {}   // empty sequence
concat {}     // empty concat
eventually {} // empty eventually
schedule {}   // empty schedule
overlap {}    // empty overlap
```

Example 226—Empty scenarios

The empty scenario always has a realization, and its realization is empty: ∅. It does not contain any action execution. Note the difference between the empty realization and empty set of realizations: the empty set of realizations means that the attempt is unsuccessful.

When an empty scenario is a member of a sequential, a concatenation, a scheduling, or an overlapping scenario, it may be dropped. For example, `sequence {{}; s;}` is equivalent to `sequence {s; {};}` and is equivalent to `s`.

A scenario is called *degenerate* if it admits an empty realization. For example, the empty scenario is degenerate because its only realization is empty. There may be non-empty degenerate scenarios, for example, `select {do a; {};}` where `a` is an action type. This scenario admits an empty realization but may also admit a realization consisting of an execution of an action of type `a`.

The top-level scenario of a cover statement shall not be degenerate. Example 227 shall result in a syntax error because the top-level scenario `select {do a; {};}` admits an empty realization. Note that cover statement `c2` is legal (`b` is a name of an action type). Though the scenario `select{ do a; {} }` is degenerate, the scenario
`sequence { select{ do a; {} }; do b; }` is not.

Copyright © 2026 Accellera. All rights reserved.

```
c1: cover {
  activity {
    select {
      do a;
      {}
    }
  }
}

c2: cover {
  activity {
    sequence {
      select {
        do a;
        {}
      }
    do b;
  }
}
```

Example 227—Degenerate scenario

16.3.8 Scheduling scenario

The *scheduling* scenario is specified by the monitor activity **schedule** statement (see Syntax 82):

```
monitor_activity_schedule_stmt::= schedule { {monitor_activity_stmt} }
```

Syntax 82—Schedule monitor activity statement

The scheduling scenario defines execution of its subscenarios in any order, provided that scenario realizations of the member scenarios are not shared. There may be any overlaps or gaps between its member scenario spans.

In the scheduling scenario, the subscenarios are matched from the checkpoint of the scheduling scenario or after it. The checkpoints of individual subscenarios are independent of each other. The realizations of scenario **schedule** { s_1,…,s_n } consist of member-wise realization unions of subscenarios s_1,. . ., and s_n, provided that these realizations of different scenarios are pairwise disjoint. For example, if set {a,b} is a realization of $s1$, set {c} is a realization of $s2$ and set {d,e,f} is a realization of s_3, then set {a,b,c,d,e,f} is a realization of **schedule** { s_1, s_2, s_3 }; here, a, b, c, d, e, and f are action executions.

Now consider monitors m defined in Example 228, the action traces shown in Figure 43, and the checkpoint t_0.

Copyright © 2026 Accellera. All rights reserved.

```
action read {}
action write {}

monitor m {
  activity {
    schedule {
      do read;
      do write;
   }
  }
}
```

Example 228—Scheduling scenario

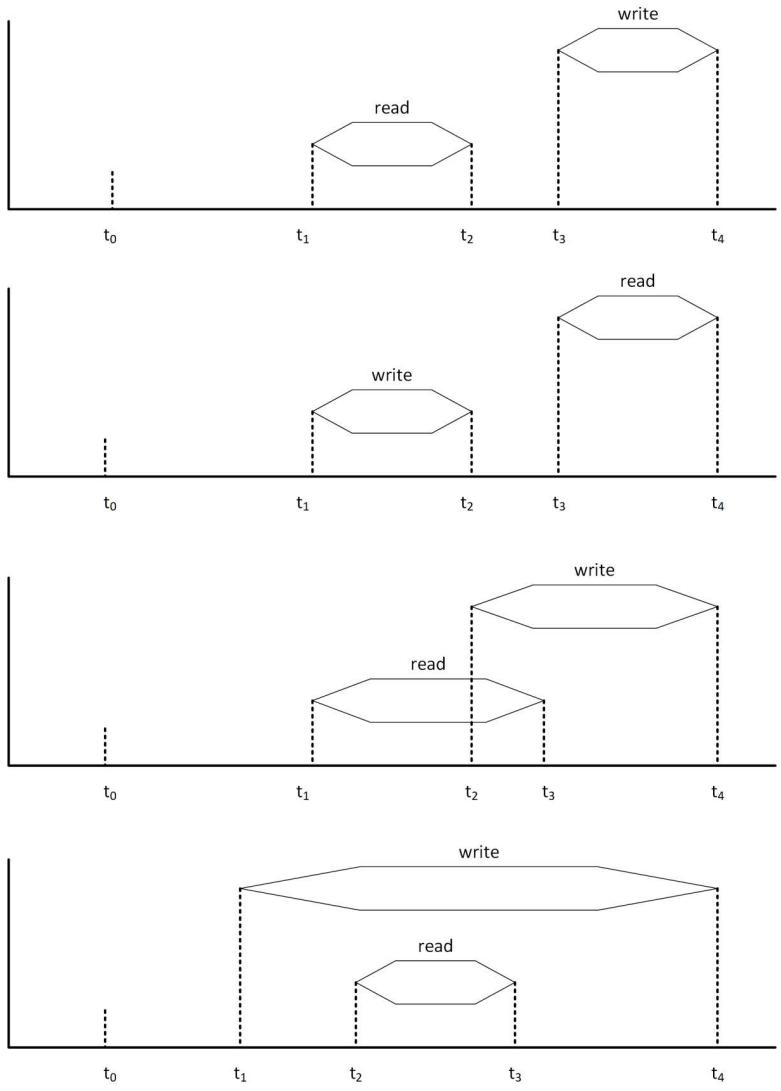


Figure 43—Scheduling scenario

Copyright © 2026 Accellera. All rights reserved.

On all traces, the only realization of the scenario defined by m is set {read, write}, its `start` time is t_1, and its match time is t_4. Note that in the first two traces, the actions are scheduled sequentially, and on the last two, they overlap.

Now consider monitor m_1 and m_2 in Example 229, the action trace shown in Figure 44, and the checkpoint t_0.

```
action read {}
action write {}
action send {}
action receive {}
monitor m1 {
  activity {
    schedule {
      sequence {
        do read;
        do write;
      };
      sequence {
        do write;
        do send;
      }
    }
  }
}

monitor m2 {
  activity {
    schedule {
      sequence {
        do write;
        do send;
      };
      sequence {
        do send;
        do receive;
      }
    }
  }
}
```

Example 229—Scheduling scenario with common actions

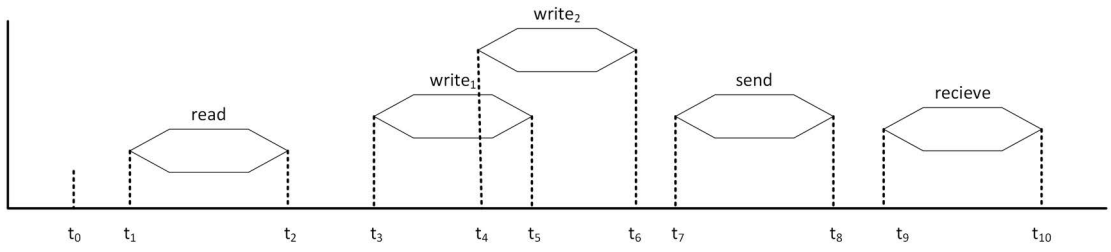


Figure 44—Scheduling scenario with common actions

The first sub-scenario of the monitor `m1` is `sequence { do read; do write; }`. It has two realizations: {read, write$_1$} and {read, write$_2$}. The second sub-scenario is

Copyright © 2026 Accellera. All rights reserved.

`sequence { do write; do send; }`, which also has two realizations: {write$_1$, send} and {write$_2$, send}. The realization of the top scenario is obtained either as a union of the first realization of the first sub-scenario and the second realization of the second sub-scenario or as a union of the second realization of the first sub-scenario and the first realization of the second sub-scenario. Both cases result in the same realization {read, write$_1$, write$_2$, send}. Other combinations of sub-scenario realizations cannot be united because they have a common element, either the first or the second `write`.

The scenario of the monitor `m2` has no match. Its first sub-scenario sequence `{ do write; do send; }` has one realization: {write$_1$, send}. Its second sub-scenario sequence `{ do send; do receive; }` also has one realization: {send, receive}. Since both realizations intersect, the top-level scenario does not have any match.

16.3.9 Monitor traversal

A monitor may be traversed within another monitor. The monitor traversal syntax is similar to the action traversal statement, see Syntax 83.

```
monitor_activity_monitor_traversal_stmt ::=
    monitor_handle_traversal_stmt
  | monitor_type_traversal_stmt

monitor_handle_traversal_stmt ::=
    monitor_identifier { [ expression ] } [ action_initializer_list ]
      monitor_activity_inline_constraints_or_empty

monitor_type_traversal_stmt ::=
  [ label_identifier : ] do monitor_type_identifier [ action_initializer_list ]
    monitor_activity_inline_constraints_or_empty
```

Syntax 83—Monitor traversal statement

At the monitor traversal statement, the scenario specified by the monitor is being matched. If the monitor traversal statement has an associated inline constraint, the monitor scenario realization shall match the specified constraint. The semantics of initializing the value of action or monitor fields via the action initializer list are the same as described in 11.3.1.2.

Example 230 illustrates a monitor traversal. Monitors `irw`, `irw1`, and `irw2` are equivalent. Consider the monitor matching in the trace shown in Figure 45 for the checkpoint t_0. The monitor `irw` specifies a sequential scenario and its only realization is {idle, read, write}. Monitor `irw1` defines a sequential scenario whose subscenarios are the traversal of action `i` with the realization {idle} and the traversal of an anonymous monitor of type `rw`. The monitor type `rw`, in its turn, defines a sequential scenario, and its checkpoint is t_1 or later so that the realization of monitor `irw1` is {idle, read, write}. The monitor `irw2` differs from the monitor `irw1` only in that it traverses the monitor of type `rw` using its handle `m`, and of course, has the same realization {idle, read, write}.

Copyright © 2026 Accellera. All rights reserved.

```
action idle {}
action read {}
action write {}
monitor rw {
  activity {
    do read;
    do write;
  }
}
monitor irw {
  idle i;
  read r;
  write w;
  activity {
    i;
    r;
    w;
  }
}
monitor irw1 {
  idle i;
  activity {
    i;
    do rw;
  }
}
monitor irw2 {
  idle i;
  rw m;
  activity {
    i;
    m;
  }
}
```

Example 230—Monitor traversal

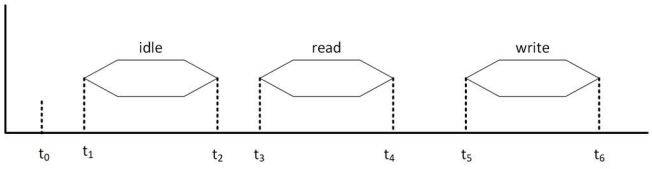


Figure 45—Monitor traversal

Example 231 illustrates using a constraint at the monitor traversal. It shows the similarities and differences between inlined and standalone constraints in monitors in the presence of sequential and concatenation scenarios.

Copyright © 2026 Accellera. All rights reserved.

```
action read {
  rand bit [32] addr;
}
action write {
  rand bit [32] addr;
}
monitor write_read_sequence {
  write w;
  read r;
  activity {
    sequence {
      w;
      r;
    }
  }
}
monitor write_read_concat {
  write w;
  read r;
  activity {
    concat {
      w;
      r;
   }
  }
}
monitor m11 {
  write_read_sequence wrs;
  activity {
    wrs with w.addr == r.addr;
  }
}
monitor m12 {
  write_read_sequence wrs;
  activity {
    wrs
  }
  constraint wrs.w.addr == wrs.r.addr;
}
monitor m21 {
  write_after_read_concat wrc;
  activity {
    wrc with w.addr == r.addr;
  }
}
monitor m22 {
  write_read_concat wrc;
  activity {
    wrc;
 }
  constraint wrc.w.addr == wrc.r.addr;
}
```

Example 231—Data constraint at monitor instantiation

Copyright © 2026 Accellera. All rights reserved.

Monitors `m11` and `m12` and monitors `m21` and `m22` are pairwise equivalent. Monitors `m11` and `m12` match when there is a `read` after a `write` from the same address. Monitors `m21` and `m22` match when the first `read` after a `write` is from the same address (see 16.3.3).

16.4 Monitor action handles and constraints

Action handles in monitors are used for readability and for constraining the monitor scenario realizations. For example, the monitor `m` in Example 232 defines a scenario capturing a `read` after a `write` from the same address.

```
monitor_activity_constraint_stmt ::= constraint monitor_constraint_set

monitor_constraint_declaration ::=
      constraint monitor_constraint_set
  | constraint identifier monitor_constraint_block

monitor_constraint_set ::=
    monitor_constraint_body_item
  | monitor_constraint_block

monitor_constraint_block ::= { { monitor_constraint_body_item } }

monitor_constraint_body_item ::=
    expression_constraint_item
  | foreach_constraint_item
  | forall_constraint_item
  | if_constraint_item
  | implication_constraint_item
  | unique_constraint_item
  | constraint_body_compile_if
  | annotation_usage monitor_constraint_body_item
  | stmt_terminator
```

Syntax 84—Monitor constraints

Copyright © 2026 Accellera. All rights reserved.

```
action read {
  rand bit [32] addr;
}
action write {
  rand bit [32] addr;
}
monitor m {
  read r;
  write w;
  activity {
    w;
    r with addr == w.addr;
  }
}
```

Example 232—Action handles in monitors

A monitor body may contain algebraic constraints (see 13.1) with the same syntax as in actions, and these constraints are subject to the same rules. As in actions, constraints in monitors may be either inline or standalone.

As explained in 16.3.1, an inline constraint imposes a condition on an action execution, see monitors `m2` and `m21` in Example 216. The standalone constraints are applied to the scenario realizations and rule out the realizations violating at least one constraint. See Example 219–Example 221 and the explanation about cover statements `c5`–`c10` there.

Inlined and standalone constraints in monitors may behave differently. A constraint inlined within an action traversal statement prescribes finding an appropriately constrained action execution of the specified type. A standalone constraint is checked at the completion of the specified statement, and if it cannot be satisfied, there is no coverage. When a monitor is built from sequences, the standalone and inlined constraints behave similarly, but when they contain a **concat** statement, the behavior may be different. This is illustrated in Example 233.

Copyright © 2026 Accellera. All rights reserved.

```
action read {
  rand bit [8] core;
}
action write {}

c1: cover {
  read r;
  write w;
  activity {
    concat {
     w;
     r with core == 1;
    }
  }
}
c2: cover {
  read r;
  write w;
  activity {
    concat {
      w;
      r;
    }
  }
  constraint r.core == 1;
}
```

Example 233—Inlined and standalone constraints in monitors

Consider the action execution trace shown in Figure 46. Cover statement `c1` is covered because its second action traversal in the concat statement looks for the next `read` action with "`core = 1`" after `write`. This action is "read$_2$". Cover statement `c2` is not covered because its second action traversal looks for the next `read` after `write` without any restrictions. This `read` action is "read$_1$", but it does not satisfy the constraint, which requires the core to be 1.

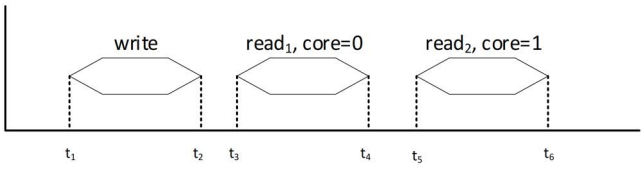


Figure 46—Inlined and standalone constraints in monitors

16.5 Covergroups in monitors

16.5.1 Covergroup sampling in monitors

Covergroups may be defined and instantiated in monitors and cover statements to collect data coverage along the scenario defined by the monitor. A monitor covergroup is sampled at the first match of the attempts of a cover statement where the monitor is traversed (directly or not). The sampling is done according to the action handle mapping associated with a first match scenario realization. If there are several first match scenario realizations, any realization may be selected for sampling by the implementation.

Copyright © 2026 Accellera. All rights reserved.

Consider the covergroup instantiation cg shown in Example 234 and the trace shown in Figure 47.

```
enum locked_e { LOCKED, UNLOCKED };
enum write_mode_e { WRITE_BACK, WRITE_THRU };

action read {
  rand locked_e lock_mode;
}
action write {
  rand write_mode_e write_mode;
}
c: cover {
  write w;
  read r;
  activity {
    w;
    r;
  }
  covergroup {
    cpw: coverpoint w.write_mode;
    cpr: coverpoint r.lock_mode;
    wXr: cross cpw, cpr;
  } cg;
}
```

Example 234—Covergroup in a cover statement

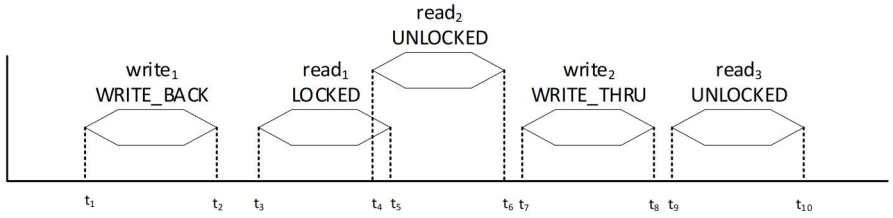


Figure 47—Covergroup in a cover statement

Successful attempts of the **cover** statement start at times t_1 and t_7. The first match scenario realization of the first attempt is {$write_1$, $read_1$} and the mapping $w \rightarrow write_1$, $r \rightarrow read_1$ so that the values WRITE_BACK and LOCKED are sampled. The first match scenario realization of the second attempt is {$write_2$, $read_3$} and the mapping $w \rightarrow write_2$, $r \rightarrow read_3$ so that the values WRITE_THRU and UNLOCKED are sampled.

Consider now the **cover** statement c in Example 235, the covergroup instantiation cg, and the trace shown in Figure 48.

Copyright © 2026 Accellera. All rights reserved.

```
enum write_mode_e { WRITE_BACK, WRITE_THRU };

action start {}
action read {}
action write {
  rand write_mode_e write_mode;
}
c: cover {
  write w;
  activity {
    do start;
    w;
    do read;
  }
  covergroup {
    cpw: coverpoint w.write_mode;
  } cg;
}
```

Example 235—Covergroup sampling for multiple realizations

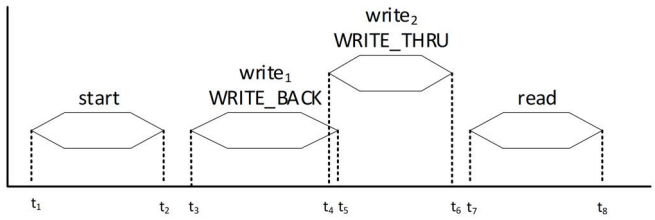


Figure 48—Covergroup sampling. Multiple realizations

There is one successful attempt of the **cover** statement `c`, starting at time t_1. It has one match time but two different realizations: {start, write$_1$, read} and {start, write$_2$, read}. In the first realization, `w` is mapped into the first `write`, and in the second one, into the second `write`. The decision whether to sample `WRITE_BACK` or `WRITE_THRU` is implementation dependent.

Example 236 shows that when the entire cover statement is not satisfied, its covergroups do not sample.

Copyright © 2026 Accellera. All rights reserved.

```
enum locked_e { LOCKED, UNLOCKED };
enum write_mode_e { WRITE_BACK, WRITE_THRU };

action read {
  rand locked_e lock_mode;
}
action write {
  rand write_mode_e write_mode;
}
action ack {}

monitor wr {
  activity {
    w: do write;
    r: do read;
  }
  covergroup {
    cpw: coverpoint w.write_mode;
    cpr: coverpoint r.lock_mode;
    wXr: cross cpw, cpr;
  } cgi;
}

c: cover {
  monitor wr;
  activity {
    wr;
    do ack;
  }
}
```

Example 236—Covergroup in a cover statement

In the trace shown in Figure 49, the covergroup data (`WRITE_BACK` and `LOCKED`) are sampled: both the **monitor** scenario and the **cover** statement top-level scenario have matches. In the trace shown in Figure 50, the covergroup data (`WRITE_BACK` and `LOCKED`) are not sampled: in spite of the fact that the **monitor** scenario has a match, the top-level **cover** statement scenario does not have a match.

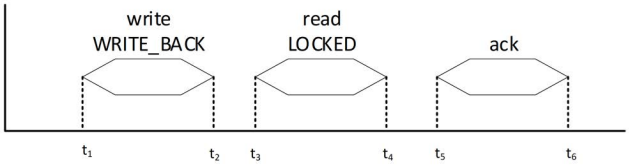


Figure 49—Covergroup instantiation in a monitor

Copyright © 2026 Accellera. All rights reserved.

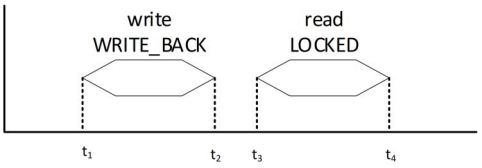


Figure 50—Covergroup instantiation in a monitor. No match of cover statement

16.5.2 Per-instance coverage in cover statements and monitors

By default, **covergroup**s collect coverage on a per-type basis (see 15.7). Per-instance coverage in **cover** statements is enabled when **per_instance** is *true* for a **covergroup** instance in a **cover** statement.

Per-instance coverage of monitors is enabled when **per_instance** is *true* for a **covergroup** instance and when there exists a contiguous path of named handles from a cover statement to the location where the covergroup is instantiated.

Copyright © 2026 Accellera. All rights reserved.

```
component pss_top {
  action write {
    rand int core;
  }
  action read {
    rand int core;
  }
  monitor mon {
    write w;
    activity { w; }
    covergroup {
      option.per_instance = true;
      cp: coverpoint w.core;
    } cg;
  }
  c1: cover {
    mon m;
    read r;
    activity {
      m;
      r with core == 0;
    }
  }
  c2: cover {
    read r;
    activity {
      do mon;
      r with core == 1;
    }
  }
  c3: cover {
    read r;
    activity {
      m: do mon;
      r with core == 2;
    }
  }
}
```

Example 237—Per-instance coverage in monitors

In Example 237, a contiguous path of named handles exists from the **cover** statement `c1` to the **covergroup** instance inside `mon`. Coverage data collected by `c1` are placed in a coverage collection unique to this named path (`m.cg`). The same is true for the **cover** statement `c3`. However, there exists no named monitor handle path from the cover statement `c2` to the **covergroup** instance inside `mon`. In this case, coverage data collected by `c2` are placed in the per-type coverage collection associated with covergroup type `mon::cg`.

16.6 Monitor activity evaluation with extension and inheritance

Monitors support both type inheritance and type extension (see Clause 17).

When a monitor inherits from another monitor, the activity declared in the inheriting monitor *shadows* (masks) the activity declared in the base monitor. The "**super;**" statement can be used to traverse the activity declared in the base monitor.

Copyright © 2026 Accellera. All rights reserved.

In Example 238, monitor `base` declares an activity that traverses action type `A`. The monitor `ext1` inherits from `base` and replaces the activity declared in `base` with an activity that traverses action type `B`. Monitor `ext2` inherits from `base` and replaces the activity declared in `base` with an activity that first traverses the activity declared in `base`, then traverses action type `C`.

```
component pss_top {
    action A { }
    action B { }
    action C { }

    monitor base {
        activity {
            do A;
        }
    }

    monitor ext1 : base {
        activity {
            do B;
        }
    }

    monitor ext2 : base {
        activity {
            super;
            do C;
        }
    }
}
```

Example 238—Monitor inheritance and traversal

Copyright © 2026 Accellera. All rights reserved.

17. Type inheritance, extension, and overrides

PSS supports the concepts of *object-oriented inheritance* and *type extension* to maximize reuse and portability of the model. *Type inheritance* allows the declaration of struct and modeling element types to be derived from a *base type* (or *supertype*), where the new *derived type* (or *subtype*) includes all attributes and other members of the base type, and allows the declaration of the derived type to add new members or mask the definition of existing members. *Type extension* allows the declaration of additional elements in an *existing* struct, modeling element type, or enumeration type, using a separate declaration.Type inheritance is described in 17.1, and type extension is described in 17.2. *Type overrides* allow type-specific and instance-specific replacement of the declared type of a field with a specified subtype, and are described in 17.5.

17.1 Type inheritance

For structs and modeling element types, the declaration may include an optional *super-spec* qualifier to declare a base type of the same type category (**action**, **monitor**, **component**, **struct**, **buffer**, **stream**, **state**, **resource**), from which the element is to be derived. The only exception is that data flow and resource objects may inherit from an element of the same type category or from a **struct**. However, if a flow or resource object type inherits from a **struct**, it only inherits the elements of the struct (fields, constraints, **exec** blocks, and so on), but it is not a struct by itself and cannot be used where a struct is expected. For example, if a flow or resource object inherits from a packed struct, it is not a packed struct by itself and cannot be used where a packed struct is expected; if a function takes a struct parameter, a flow or resource object that inherits from this struct cannot be passed for this parameter in a function call, and so on.

A derived type includes all elements from the base type, and may declare new elements that may or may not have the same name as a corresponding element in the base type. For fields declared in a derived type with the same name as a field in the base type, the derived type's field shadows (masks) the base type's field, and the base type's field may be referenced as "**`super.`**`<name>`". Certain unnamed elements, such as *activities* and procedural *exec* blocks, may invoke the corresponding element(s) from the base type by the "**`super;`**" statement.

The behavior of specific elements when declared in a derived type is shown in Table 27.

Copyright © 2026 Accellera. All rights reserved.

Table 27—Derived type element behaviors

Element kind	In a component	In an action	In a monitor	In a struct, data flow or resource object
type	shadow	n/a	shadow	n/a
action (override)	Override. Use "super" to refer to elements in base type.	n/a	n/a	n/a
activity	n/a	shadow, may call **super;**	shadow, may call **super;**	n/a
named sub-activities	n/a	shadow (may access base activity state-ment as **super.**name)	shadow (may access base activity state-ment as **super.**name)	n/a
generic constraint	n/a	shadow (may access base constraint as **super.**name)	n/a	shadow (may access base constraint as **super.**name)
named constraint	n/a	shadow	shadow	shadow
unnamed constraint	n/a	added	added	added
field	shadow (may access base field as **super.**name)	shadow (may access base field as **super.**name)	shadow (may access base field as **super.**name	shadow (may access base field as **super.**name)[a]
instance function	shadow (may call base function as **super.**name(args))	n/a	n/a	n/a
static function	shadow	n/a	n/a	n/a
override declaration	added	added	added	n/a
object pool bind	added	n/a	n/a	n/a
procedural exec block	shadow, may call **super;**	shadow, may call **super;**	n/a	shadow, may call **super;**
target-template exec block	shadow	shadow	n/a	shadow

[a] If field is not a pool instance. Accessing the pool instance of a supertype component to do a bind in the subtype is not allowed.

Activities in derived actions and monitors shadow the activities from the base action or monitor type. However, the "**`super;`**" statement may be used to traverse the base activity (or activities). See Example 119 in 11.6 and Example 238 in 16.6.

Procedural *exec* blocks defined in a derived type shadow same-kind *exec* block(s) defined in the base type. The *exec* block in the derived type may include the "**`super;`**" statement, which will execute the contents of the corresponding base-type *exec* block(s) at that point. See 20.1.4.1 and 20.1.4.2.

Copyright © 2026 Accellera. All rights reserved.

Target-template *exec* blocks defined in a derived type shadow same-kind *exec* blocks with the same target language identifier in the base type. The "**`super;`**" statement shall not be allowed in a target-template *exec* block.

Example 239 shows a simple case of declaring a component `base_c`, which contains an action declaration, `base_a`. Derived component `der_c` inherits from `base_c`, so it is treated as having action `base_a` already declared within it. Note that `base_c` and `der_c` are different component types. Action `der_a` inherits from `base_a`, so it already includes random integer `i` and bit-vector `b`, as well as the unnamed constraint limiting `i` to be less than `10` and constraint `c` forcing `b > 7`. Derived action `der_a` adds an additional random integer, `j`, a new unnamed constraint that relates the values of `i` and `j`, and a new constraint `c` that relates the values of `b` and `j`, shadowing constraint `c` from action `base_a`.

```
component base_c {
  action base_a {
    rand int i;
    rand bit[32] b;
    constraint {i < 10;}
    constraint c {b > 7;}
  }
}

component der_c : base_c {
  action der_a : base_a {
    rand int j;
    constraint {j > 5 -> i < 5;}
    constraint c {j < 10 -> b < 128;}
  }
}
```

Example 239—Declaring derived components and actions

When a pool **bind** statement (see 12.3) is used in a base component type, it may also apply to a derived type, provided that any new component instances and actions in the derived type also match the path specification in the **bind** statement and that the types of the object references match the pool type exactly.

In Example 240, the default **bind** statement in `base_c` binds the `cpu_p` pool to the actions `act1_a` and `act2_a` defined therein. Since `der_c` is derived from `base_c`, it also inherits the **bind** statement, which applies to all action definitions in `der_c` that match the path specification. In the context of `der_c`, the default bind statement binds all three actions `act1_a`, `act2_a` and `act3_a` to the `cpu_p` pool.

Copyright © 2026 Accellera. All rights reserved.

```
resource cpu_core_s {...
}

component base_c {
  pool[4] cpu_core_s cpu_p;
  bind cpu_p *;
  action act1_a {
    share cpu_core_s cpu_share;
  }
  action act2_a {
    lock cpu_core_s cpu_lock;
  }
}

component der_c : base_c {
  action act3_a {
    share cpu_core_s cpu_share;
  }
  ...
}
```

Example 240—Default pool with inheritance

As mentioned above, a derived type inherits all members from the base type and may declare additional elements specific to the derived type. When a named element is declared in the derived type with the same name as an element in the base type, the derived type's declaration shadows (masks) the base type's declaration (as with constraint `c` in Example 239).

When the shadowed element is a function, the function call is *polymorphic*, that is, the actual function called depends on its context component. In Example 241, component `der_c` shadows the definition of function `foo()` in component `base_c`. Action `call_foo` invokes the appropriate definition of `foo()` depending on the type of its context component. Action `test` schedules `call_foo` in the context of a component of type `base_c`, followed by `call_foo` in the context of `der_c`. Executing `test` will call the core library target function **`message()`** to add the following messages to the execution log, at **`LOW`** verbosity:

```
base_c::foo
der_c::foo
```

Copyright © 2026 Accellera. All rights reserved.

```
import std_pkg::*;

component base_c {
   target function void foo() {
      message(LOW, "base_c::foo");
   }

   action call_foo {
      exec body {
         comp.foo();
      }
   }
}

component der_c : base_c {
   function void foo() {
      message(LOW, "der_c::foo");
   }
};

component pss_top {
   base_c b;
   der_c d;
   action test {
      base_c::call_foo b_foo, d_foo;
      constraint {b_foo.comp == this.comp.b;
                  d_foo.comp == this.comp.d;}
      activity {
         b_foo;
         d_foo;
      }
   }
}
```

Example 241—Polymorphic function calls

As discussed in 9.1.3, the qualified name of an action declared in a component is of the form '*component-type*::*action-type*'. In Example 242, the base component `dma_base_c` declares action `xfer_a`. The derived component `dma_der_c` declares the compound action `mult_xfer_a`, which traverses the `xfer_a` action. Since `dma_der_c` inherits the `xfer_a` action, the anonymous (by type) traversal in `mult_xfer_a` correctly resolves to the `xfer_a` action declared in the base component. It is thus not necessary to further qualify the type name `xfer_a` in the anonymous traversal in `mult_xfer_a`.

The component `dma_test_c` instantiates the derived component `dma_der_c`. The first traversal statement in the activity is an anonymous traversal of the `dma_der_c::mult_xfer_a` action. The next statement anonymously traverses the `dma_base_c::xfer_a` action. We can use the `dma_base_c` path qualifier because the instantiated subcomponent of type `dma_der_c` *is also* considered a `dma_base_c` component. It would be illegal to refer to `dma_base_c::mult_xfer_a` because `mult_xfer_a` is not declared in `dma_base_c`. To promote reuse, the third anonymous traversal statement is preferred, referring to `dma_der_c::xfer_a`, since `xfer_a` can be used without knowing whether it was declared in the base component or the derived component. Note that, since there is only a single instance of the `dma_der_c` component, the instance context of these traversals is the same.

Copyright © 2026 Accellera. All rights reserved.

```
component dma_base_c {
  action xfer_a {
    ...
  }
}

component dma_der_c : dma_base_c {
  action mult_xfer_a {
    activity {
      repeat(3) {
        do xfer_a; // dma_base_c::xfer_a
      }
    }
  }
}

component dma_test_c {
  dma_der_c dma;

  action test_a {
    activity {
      do dma_der_c::mult_xfer_a;
      do dma_base_c::xfer_a;
      do dma_der_c::xfer_a; // dma_base_c::xfer_a
    }
  }
}
```

Example 242—Derived type is also a base type

In Example 243, there are two instances of the `dma_der_c` component instantiated in `dma_test_c`. For the first anonymous traversal of `dma_base_c::xfer_a`, either instance may be chosen as context for the `xfer_a` action. In the second anonymous traversal, the **`comp`** attribute is constrained to specify that the context component shall be `dma_test_c.dma1`. As stated in 9.1.5, the static type of the **`comp`** attribute of `dma_der_c::xfer_a` is actually `dma_base_c`, since that is its containing component type (See also 13.1.4).

Because **`comp`** is of type `dma_base_c` and not `dma_der_c`, it would be illegal to refer to fields of `dma_der_c` as relative to **`comp`**, since these fields are not in `dma_base_c`. Rather, fields of `dma_der_c` may be referred to relative to `this.comp.dma1`, which is the actual instance of `dma_der_c` (which is also a `dma_base_c`) in which `xfer_a` will execute. Thus, based on the actual instance of a context component, we can constrain the fields of `xfer_a` even though `xfer_a` may not have visibility otherwise to the `dma_der_c` fields that control the constraints.

Copyright © 2026 Accellera. All rights reserved.

```
component dma_base_c {
  action xfer_a {
    rand int i;
    ...
  }
}

component dma_der_c : dma_base_c {
  int j;
  action mult_xfer_a {
    activity {
      repeat(3) {
        do xfer_a; // dma_base_c::xfer_a
      }
    }
  }
}

component dma_test_c {
  dma_der_c dma1, dma2;

  action test_a {
    activity {
      do dma_base_c::xfer_a;
      do dma_der_c::xfer_a with {comp == this.comp.dma1;
                                 (this.comp.dma1.j < 8) -> i>4;};
    }
  }
}
```

Example 243—Use of comp and this.comp with inheritance

When declaring a new component, it shall be illegal to declare types that derive from types declared in an existing component type unless the new component derives from the existing component.

Example 244 demonstrates why this kind of inheritance is problematic. Action `new_a`, derived from `existing_c::existing_a`, inherits constraint `con` that constrains `k` based on the value of attribute `i` of component `existing_c`. The **`comp`** field of action `new_a` is of type `new_c` and not `existing_c`, and therefore does not have attribute `i`. For that reason, the action `new_a` is not able to evaluate constraint `con`, Thus, modeling with this kind of inheritance cannot work.

Copyright © 2026 Accellera. All rights reserved.

```
component existing_c {
  int i;
  exec init_up {i = 1;}
  action existing_a {
    rand int in [0..4] k;
    constraint con {k > comp.i;};
  }
}

component new_c {
  action new_a : existing_c::existing_a {} // Illegal
}

component pss_top {
  new_c c;
  action entry_a {
    activity {
      do new_c::new_a;
    }
  }
}
```

Example 244—Illegal inheritance declaration

17.2 Type extension

*Type extension*s in PSS enable the decomposition of model code so as to maximize reuse and portability. Modeling elements and data types may have a number of properties that are logically independent. Moreover, distinct concerns with respect to the same entities often need to be developed independently. Later, the relevant definitions need to be integrated, or woven into one model, for the purpose of generating tests.

Some typical examples of concerns that cut across multiple model entities are:

- Implementation of actions and objects for, or in the context of, some specific target platform/language.
- Model configuration of generic definitions for a specific device under test (DUT) / environment configuration, affecting components and data types that are declared and instantiated elsewhere.
- Definition of functional elements of a system that introduce new properties to common objects, which define their inputs and outputs.
- Restricting monitors with additional constraints.

Such crosscutting concerns can be decoupled from one another by using type extensions and then encapsulated as **packages** (see 18.1).

Modeling element, **struct**, and enumeration types in PSS are extensible. They are declared once, along with their initial definition, and may later be extended any number of times, with new body items being introduced into their scope. Items introduced in extensions may be of the same kinds as those introduced in the initial definition. Extension statements may appear in **package** and **component** definitions.

An extension statement explicitly specifies the kind of type being extended, which shall agree with the specific type named (see Syntax 85).

Copyright © 2026 Accellera. All rights reserved.

The overall definition of any given type in a model is the sum total of its definition statements—the initial one along with extensions in active packages (see 18.1). The semantics of extensions are those of weaving all those statements into a single definition.

Every type extension, regardless of whether it extends a package-level type or a component-level inner type, is associated with the nearest **package** that lexically encloses its definition (an explicit **package** if enclosed in a *package_declaration* statement or otherwise the unnamed global package (see 18.1)).

Members introduced in an extension of a type can be referenced throughout the package in which they were introduced. As a corollary, members introduced in extensions associated with the global package can be referenced everywhere. Members introduced in extensions cannot be referenced *outside* the scope of the package in which the extension is defined unless the reference occurs in a lexical scope that wildcard-imports that package.

These rules concern reference of static members as well as non-static members, and apply regardless of whether fully-qualified static paths are used (for static members).

17.2.1 Syntax

```
extend_stmt ::=
    extend action type_identifier { { action_body_item } }
  | extend monitor type_identifier { { monitor_body_item } }
  | extend component type_identifier { { component_body_item } }
  | extend struct_kind type_identifier { { struct_body_item } }
  | extend enum type_identifier { [ enum_item { , enum_item } ] }
  | extend annotation type_identifier { { annotation_body_item } }
```

Syntax 85—Type extension

17.2.2 Examples

Examples of type extension are shown in Example 245 and Example 246.

```
enum config_modes_e {UNKNOWN, MODE_A=10, MODE_B=20};

component uart_c {
    action configure {
        rand config_modes_e mode;
        constraint {mode != UNKNOWN;}
    }
}

package additional_config_pkg {
    extend enum config_modes_e {MODE_C=30, MODE_D=50}

    extend action uart_c::configure {
        constraint {mode != MODE_D;}
    }
}
```

Example 245—Type extension

Copyright © 2026 Accellera. All rights reserved.

```
component pcie_c {
    action read { bit [32] addr; }
    action write { bit [32] addr; }

    monitor read_after_write {
        read r;
        write w;
        activity {
            w;
            r;
       }
    }
}

package additional_checks_pkg {

    extend monitor pci_c::read_after_write {
        constraint { r.addr == w.addr; }
    }
}
```

Example 246—Monitor type extension

17.2.3 Struct and modeling element type extensions

Any kind of member declared in the context of the initial definition of a **struct** or modeling element type can be declared in the context of an extension, as per its entity category (**action**, **monitor**, **component**, **buffer**, **stream**, **state**, **resource**, or **struct**).

Named type members of any kind, fields in particular, may be introduced in the context of a type extension. Names of fields introduced in an extension shall not conflict with those declared in the initial definition of the type. They shall also be unique in the scope of their type within the **package** in which they are declared. However, field names do not have to be unique across extensions of the same type in different packages.

Fields are always accessible within the scope of the package in which they are declared, shadowing (masking) fields with the same name declared in other packages. Members declared in a different package are accessible if the declaring package is wildcard-imported into the scope of the accessing **package** or **component**, given that the reference is unique. If the same field name or type name is wildcard-imported from two or more separate packages, it shall be an error to reference it.

In Example 247, an **action** type is initially defined in the context of a **component** and later extended in a separate **package**. Ultimately the **action** type is used in a compound action of a parent **component**. The **component** explicitly wildcard-imports the **package** with the extension and can therefore constrain the attribute introduced in the extension.

Copyright © 2026 Accellera. All rights reserved.

```
component mem_ops_c {
    enum mem_block_tag_e {SYS_MEM, A_MEM, B_MEM, DDR};

    buffer mem_buff_s {
        rand mem_block_tag_e mem_block;
    }

    pool mem_buff_s mem;
    bind mem *;

    action memcpy {
        input mem_buff_s src_buff;
        output mem_buff_s dst_buff;
    }
}

package soc_config_pkg {
    extend action mem_ops_c::memcpy {
        rand int in [1, 2, 4, 8] ta_width; // introducing new attribute

        constraint { // layering additional constraint
            src_buff.mem_block in [SYS_MEM, A_MEM, DDR];
            dst_buff.mem_block in [SYS_MEM, A_MEM, DDR];
            ta_width < 4 -> dst_buff.mem_block != A_MEM;
        }
    }
}

component pss_top {
    import soc_config_pkg::*;// explicitly importing the package grants
                             // access to types and type members
    mem_ops_c mem_ops;

    action test {
        mem_ops_c::memcpy cpy1, cpy2;
        constraint cpy1.ta_width == cpy2.ta_width;// constraining an
                                 // attribute introduced in an extension
        activity {
            repeat (3) {
                parallel { cpy1; cpy2; };
            }
        }
    }
}
```

Example 247—Action type extension

17.2.4 Enumeration type extensions

Enumeration types can be extended in one or more package contexts, introducing new enum items to the domain of all variables of that type. Each enum item in an **enum** type shall be associated with an integer value that is unique across the initial definition and all the extensions of the type. Enum item values are assigned according to the same rules they would be if all the enum items appeared in the initial definition, according to the order of package evaluations. An explicit conflicting value assignment shall be illegal.

Copyright © 2026 Accellera. All rights reserved.

An enum item introduced in an extension can be referenced within the **package** in which the extension is defined. Outside that **package**, enum items can be referenced inside a lexical scope that wildcard-imports the respective package.

In Example 248, an **enum** type is initially declared empty and later extended in two independent **package**s. Ultimately items are referenced from a **component** that wildcard-imports both **package**s.

```
package mem_defs_pkg { // reusable definitions
    enum mem_block_tag_e {}; // initially empty

    buffer mem_buff_s {
        rand mem_block_tag_e mem_block;
    }
}
package AB_subsystem_pkg {
    import mem_defs_pkg ::*;

    extend enum mem_block_tag_e {A_MEM, B_MEM};
}
package soc_config_pkg {
    import mem_defs_pkg ::*;

    extend enum mem_block_tag_e {SYS_MEM, DDR};
}
component dma_c {
  import mem_defs_pkg::*;
  action mem2mem_xfer {
    input mem_buff_s src_buff;
    output mem_buff_s dst_buff;
 }
}
extend component dma_c {
    import AB_subsystem_pkg::*; // wildcard-importing the package
    import soc_config_pkg::*;   // grants access to enum items

    action dma_test {

        activity {
            do mem2mem_xfer with {
                src_buff.mem_block == A_MEM;
                dst_buff.mem_block == DDR;
            };
        }
    }
}
```

Example 248—Enum type extensions

17.2.5 Ordering of type extensions

Multiple type extensions of the same type can be coded independently, and be integrated and woven into a single stimulus model, without interfering with or affecting the operation of one another. Methodology should encourage making no assumptions on their relative order.

Copyright © 2026 Accellera. All rights reserved.

From a semantics point of view, order would be visible in the following cases:

- Invocation order of *exec block*s of the same kind
- Multiple **default** value constraints, **default disable** constraints, **soft** constraints, and type override declarations occurring in a scope of the same type
- Integer values associated with enum items that do not explicitly have a value assignment

The initial definition always comes first in ordering of members. The order of extensions conforms to the order in which packages are processed by a PSS implementation.

NOTE—This standard does not define specific ways in which a user can control the package processing order.

17.2.6 Template type extensions

Template types, as all other user-defined types, may be extended using the **extend** statement.

Template types may be extended in two ways:

a) Extending the generic template type. The extension will apply to all instances of the template type.
b) Extending the template type instance. The extension will apply to all instances of the template type that are instantiated with the same set of parameter values.

NOTE—Partial template specialization is not supported.

17.2.6.1 Examples

Examples of extending the generic template type and the template type instance are shown in Example 249.

Copyright © 2026 Accellera. All rights reserved.

```
struct domain_s <int LB = 4, int UB = 7> {
  rand int attr;
  constraint attr >= LB && attr <= UB;
}

struct container_s {
  domain_s<2, 7> domA;          // specialized with LB = 2, UB = 7
  domain_s<2, 8> domB;          // specialized with LB = 2, UB = 8
}

extend struct domain_s {
  rand int attr_all;            // container_s::domA and container_s::domB
                                // will have attr_all
  constraint attr_all > LB && attr_all < UB;
}

extend struct domain_s<2> {     // extend instance specialized with
                                // LB = 2, UB = 7 (default)
  rand int attr_2_7;            // container_t::domA will have attr_2_7
  constraint attr_2_7 > LB && attr_2_7 < UB; // parameters accessible in
                                             // template instance extension
}

struct sub_domain_s<int MIN, int MAX> : domain_s<MIN, MAX> {
  rand int domain_size;
  constraint domain_size == MAX - MIN + 1;

constraint half_max_domain() {
    attr >= LB && attr <= UB/2; // Error - LB and UB parameters not accessible
                                                    // in inherited struct
  }
}
```

Example 249—Template type extension

In the example above, the generic template type extension is used to add `attr_all` to all instances of `domain_s`. The template type instance extension is used to add `attr_2_7` to the specific `<2,7>` instance of `domain_s`.

17.3 Combining inheritance and extension

It is important to understand that *inheritance* creates a *new* type derived from the base type, while *extension modifies* the definition of an *existing* type. Once a derived type is created by inheriting from a base type, the derived type may be extended just as any other type. In this case, the extensions to the derived type do not affect the base type. However, since a derived type inherits from its base type, any extensions to the base type will also affect the derived type. If multiple types are derived from the same base type, extensions to the base type will affect all derivations thereof.

Extending types in a component scope is only allowed for types that are defined in that scope. It shall be illegal to extend a type defined in a base component type from a derived or unrelated component type.

Copyright © 2026 Accellera. All rights reserved.

In Example 250, by extending action `der_a` in component `der_c`, we add a new constraint on the `j` field. This constraint is added to the existing constraints in the initial definition of `der_a`. By extending action `base_a` in the `base_c` extension, we add a new constraint, `i > 2`, which is then inherited by the derived action, `der_a`. The result is that `j` is constrained to be greater than `7`, implying that `i` shall be less than `5`, and the additional constraint requires that `i` shall also be greater than `2`.

The attempt to extend action `base_a` in component `der_c` is illegal, since `base_a` was originally declared in `base_c`, which is a different type from `der_c`.

```
component base_c {
  action base_a {
    rand int i;
    rand bit[32] b;
    constraint { i < 10; }
    constraint c { b > 7; }
  }
}

component der_c : base_c {
  action der_a : base_a {
    rand int j;
    constraint { j > 5 -> i < 5; }
    constraint c { j < 10 -> b < 128; }
  }

  extend action der_a {
    constraint { j > 7; }
  }

  extend action base_a {...} // ILLEGAL
}

extend component base_c {
  extend action base_a {
    constraint { i > 2; }
  }
}
```

Example 250—Combining inheritance and extension

In Example 251, in the **`pss_top`** root action, the anonymous traversal of `der_c::base_a` will use the `base_a` action as extended in `base_c` in the global scope. Thus, the constraints `i > 2` and `i < 10` will apply. Its execution context will be either instance `c1` or `c2` of `der_c`.

The anonymous traversal of `der_c::der_a` similarly will use the extended definition of `der_a`, but the **`with`** constraint forces the execution context to be instance `c1`. Note that the constraint `c` in `der_c::der_a` masks the original constraint `c` in `base_c::base_a`, so the resolved set of applicable constraints will be:

- `j > 7`
- `i < 5` (due to constraint `j > 5 -> i < 5`)
- `j < 10 -> b < 128`

Copyright © 2026 Accellera. All rights reserved.

```
component base_c {
  action base_a {
    rand int i;
    rand bit[32] b;
    constraint { i < 10; }
    constraint c { b > 7; }
  }
}

component der_c : base_c {
  action der_a : base_a {
    rand int j;
    constraint { j > 5 -> i < 5; }
    constraint c { j < 10 -> b < 128; }
  }
}

extend component der_c {
  extend action der_a {
    constraint { j > 7; }
  }
}

extend component base_c {
  extend action base_a {
    constraint { i > 2; }
  }
}

component pss_top {
  der_c c1, c2;

  action root {
    activity {
      do der_c::base_a;
      do der_c::der_a with {comp == this.comp.c1; };
    }
  }
}
```

Example 251—Inheritance and extension of constraints

17.4 Access protection

By default, all data attributes of **struct**s and model element types have public accessibility. The default accessibility can be modified for a single data attribute by prefixing the attribute declaration with the desired accessibility. The default accessibility can be modified for all attributes going forward by specifying a block-access modifier.

The following also apply:

a) A **public** attribute is accessible from any element in the model.

b) A **private** attribute is accessible only from the element in which the attribute is declared.

c) A **protected** attribute is accessible only from the element in which the attribute is declared, from sub-elements that inherit from it, and from their extensions.

Copyright © 2026 Accellera. All rights reserved.

Example 252 shows using a per-attribute access modifier to change the accessibility of the random attribute `b`. Fields `a` and `c` are publicly accessible.

```
struct S1 {
    rand int a;         // public accessibility (default)
    private rand int b; // private accessibility
    rand int c;         // public accessibility (default)
}
```

Example 252—Per-attribute access modifier

Example 253 shows using block access modifiers to set the accessibility of a group of attributes. Fields `w` and `x` are private due to the **private:** directive. Field `y` is public because its access modifier is explicitly specified. Field `z` is private, since the **private:** block access modifier is in effect. Field `s` is public, since the preceding **public:** directive has changed the default accessibility back to public.

```
struct S2 {
    private:
        rand int w;         // private accessibility
        rand int x;         // private accessibility
        public rand int y;  // public accessibility
        rand int z;         // private accessibility

    public:
        rand int s;         // public accessibility
}
```

Example 253—Block access modifier

17.5 Overriding types

The **override** block (see Syntax 86) allows type- and instance-specific replacement of the declared type of a field with some specified subtype.

Overrides apply to action and monitor fields, struct attribute fields, and component instance fields. In the presence of **override** blocks in the model, the actual type that is instantiated under a field is determined according to the following rules:

a) Walking from the field up the hierarchy from the contained entity to the containing entity, the applicable **override** directive is the one highest up in the containment tree.

b) Within the same container, **instance** override takes precedence over **type** override.

c) For the same container and kind, an override introduced later in the code takes precedence.

Overrides do not apply to reference fields, namely fields with the modifiers **input**, **output**, **lock**, and **share**. Component-type overrides under actions and monitors as well as action-type and monitor-type overrides under components are not applicable to any fields; this shall be an error.

Copyright © 2026 Accellera. All rights reserved.

17.5.1 Syntax

```
override_declaration ::= override { { override_stmt } }
override_stmt ::=
    type_override
  | instance_override
  | override_compile_if
  | annotation_usage override_stmt
  | stmt_terminator
type_override ::= type type_identifier with type_identifier ;
instance_override ::= instance hierarchical_id with type_identifier ;
```

Syntax 86—Override declaration

17.5.2 Example

Example 254 combines type- and instance-specific overrides with type inheritance. Action `reg2axi_top` specifies that all `axi_write_action` instances shall be instances of `axi_write_action_x`. The specific instance `xlator.axi_action` shall be an instance of `axi_write_action_x2`. Action `reg2axi_top_x` specifies that all instances of `axi_write_action` shall be instances of `axi_write_action_x4`, which supersedes the override in `reg2axi_top`. In addition, action `reg2axi_top_x` specifies that the specific instance `xlator.axi_action` shall be an instance of `axi_write_action_x3`.

Copyright © 2026 Accellera. All rights reserved.

```
action axi_write_action { ... };

action xlator_action {
  axi_write_action axi_action;
  axi_write_action other_axi_action;
  activity {
    axi_action; // overridden by instance
    other_axi_action; // overridden by type
  }
};

action axi_write_action_x : axi_write_action { ... };

action axi_write_action_x2 : axi_write_action_x { ... };

action axi_write_action_x3 : axi_write_action_x { ... };

action axi_write_action_x4 : axi_write_action_x { ... };

action reg2axi_top {
  override {
    type axi_write_action with axi_write_action_x;
    instance xlator.axi_action with axi_write_action_x2;
  }

  xlator_action   xlator;
  activity {
    repeat (10) {
      xlator; // override applies equally to all 10 traversals
    }
  }
};
action reg2axi_top_x : reg2axi_top {
  override {
    type axi_write_action with axi_write_action_x4;
    instance xlator.axi_action with axi_write_action_x3;
  }
};
```

Example 254—Type inheritance and overrides

Copyright © 2026 Accellera. All rights reserved.

18. Source organization and processing

A PSS model is captured in one or more *source units*. Source units contain declarations of PSS elements. Name resolution rules for types are specified with respect to source units. The bounds of a source unit are specified either by a single file or by a collection of files identified to the PSS processing tool as being part of a single source unit. The files comprising a multi-file source unit could be identified to the PSS processing tool in several different ways. For example, the PSS processing tool could be instructed to consider all PSS source files in a given directory to be a single source unit. The PSS processing tool could be instructed to consider all PSS source files listed in a filelist to be a single source unit. Tool implementations shall support both single-file and multi-file source unit processing modes, but this standard does not dictate the mechanism by which source units shall be specified to the PSS processing tool.

A lexical scope shall be fully contained within a single source file, independent of whether source files are processed as single- or multi-file source units.

The processing order of a set of source units is user-specified to the PSS processing tool. This standard does not dictate a specific processing order for files *within* a multi-file source unit, but tools may provide users with means to control it.

18.1 Packages

Packages are a way to group, encapsulate, and identify sets of related definitions, namely type declarations and type extensions. In a verification project, some definitions may be required for the purpose of generating certain tests, while others need to be used for different tests. Moreover, extensions to the same types may be inconsistent with one another, e.g., by introducing contradicting constraints or specifying different mappings to the target platform. By enclosing these definitions in packages, they may coexist and be managed more easily.

Packages also constitute namespaces for the types, functions, and constants declared in their scope. From a namespace point of view, **package**s and **component**s have the same meaning and use (see also 9.1.3). However, in contrast to **component**s, **package**s cannot be instantiated, and cannot contain attributes, sub-component instances, or concrete **action** definitions.

Type declarations, functions, and constants declared under the scope of a **package** declaration statement are members of that package. Package members may be referenced from outside the package using a qualified reference or made visible by *importing* them into the referencing scope (see 18.1.3).

Definition statements that do not occur inside the lexical scope of a **package** declaration are implicitly associated with the *unnamed global package*. Elements in the unnamed global package are visible to all user-defined namespaces without the need for an **import** statement.

Tools may provide means to control and query which packages are active in the generation of a given test. Tools may also provide ways to locate source files of a given package in the file system. However, these means are not covered herein.

Copyright © 2026 Accellera. All rights reserved.

18.1.1 Package declarations

18.1.1.1 Syntax

```
package_declaration ::= package package_id_path { { package_body_item } }
package_id_path ::= package_identifier { :: package_identifier }
package_identifier ::= identifier
package_body_item ::=
    abstract_action_declaration
  | struct_declaration
  | enum_declaration
  | covergroup_declaration
  | function_decl
  | import_class_decl
  | procedural_function
  | import_function
  | target_template_function
  | export_action
  | typedef_declaration
  | import_stmt
  | extend_stmt
  | const_field_declaration
  | component_declaration
  | abstract_monitor_declaration
  | package_declaration
  | compile_assert_stmt
  | package_body_compile_if
  | generic_constraint_declaration
  | annotation_declaration
  | annotation_usage package_body_item
  | export_function
  | stmt_terminator

const_field_declaration ::= [ static ] const data_declaration
```

Syntax 87—Package declaration

The following also apply:

a) Multiple **package** statements can apply to the same package name. The **package** contains the members and type extensions declared in all package scopes with the same name.

b) In a *const_field_declaration*, the **static** keyword is optional, but the field is a static constant even if the **static** keyword is not used.

Copyright © 2026 Accellera. All rights reserved.

18.1.1.2 Examples

For an example of package usage, see 20.2.7.

18.1.2 Nested packages

A package may be *nested* inside another package. There are two way to declare a *nested package*.

One way is to include a package declaration inside the outer package declaration, as shown in the following example:

```
package my_lib {
  package impl {
    struct internal_impl_s {}
  }
}
```

Example 255—Hierarchical declaration of nested package

In the example above, the fully-qualified type name of the **struct** `internal_impl_s` is `my_lib::impl::internal_impl_s`.

Nested packages can also be specified with double-colon-separated package identifier paths. In the example below, the fully-qualified type name of the **struct** `internal_impl_s` is also `my_lib::impl::internal_impl_s`.

```
package my_lib::impl {
  struct internal_impl_s {}
}
```

Example 256—Direct declaration of nested package

Declaring a package inside another is equivalent to directly specifying a hierarchical name for a package namespace

The declaration order of package namespaces is not significant. So, for example, it is not necessary to declare an outer namespace prior to declaring an inner namespace. In the example below, two **structs** are declared. `my_lib::impl::internal_impl_s` is declared first, while `my_lib::public_s` is declared second.

```
package my_lib::impl {
   struct internal_impl_s {}
}

package my_lib {
   struct public_s {}
}
```

Example 257—Declaration of nested package before outer package

Copyright © 2026 Accellera. All rights reserved.

18.1.3 Referencing package members

There are three ways to reference package members from outside the scope of their declaring package: *qualified reference*, *explicit import*, and *wildcard import*.

One way to use a declaration from a package is to reference it explicitly using the scope resolution operator **::**. This is called a *qualified reference*. Example:

```
my_lib::public_s my_struct;
```

An alternate method for referencing package declarations is via the **import** statement. Importing an identifier into a **package** or **component** makes that identifier visible within that lexical scope without requiring the scope resolution operator. An **import** statement is a name resolution directive, and does not introduce symbol declarations or symbol aliases into the namespace in which it appears.

Two forms of the **import** statement are provided: *explicit import* and *wildcard import*. An *explicit import* only imports the symbols specifically referenced by the **import**. Example:

```
import my_lib::public_s;
public_s my_struct;
```

It shall be illegal to explicitly import an identifier from a package if the same name is already declared in the importing namespace or to explicitly import the same identifier from two different **package**s.

A *wildcard import* allows all identifiers declared within a package to be imported into a lexical scope, provided the identifier is not otherwise defined anywhere in the importing **component** or **package**. A wildcard import also allows access from the lexical scope to members declared in type extensions found in the imported package. Note that type extensions are unnamed and therefore cannot be explicitly imported.

A wildcard import is of the following form:

```
import my_lib::*;
public_s my_struct;
```

A local declaration of an identifier takes precedence over a wildcard import of the same identifier. An explicit import of an identifier takes precedence over a wildcard import of the same identifier from a different package. If the same name is declared in two wildcard-imported packages, neither is imported, a qualified reference shall be used.

import specifications may only appear in one of the following lexical scopes:

- In the global scope within a source file
- In a **package** declaration statement
- In a **component** declaration or extension statement

In all these cases, **import** specifications may be conditional within a **compile if** statement (see 19.2).

In all cases, **import** specifications shall appear first in any of these lexical scopes. In other words, no statement other than another **import** or a **compile if** statement shall precede an **import** statement within the lexical scope.

The effect of an **import** statement shall be limited to the lexical scope in which it appears, including any sub-scopes, namely:

Copyright © 2026 Accellera. All rights reserved.

— The effect of **import** in the global scope is limited to the same source file. If the source file is part of a multi-file source unit, the **import** shall *not* have effect in other files of the same source unit.
— The effect of **import** in a **package** or **component** declaration or a **component** extension is limited to this declaration or extension. It shall *not* have effect in other declarations or extensions of the same package or component.

Elements in the unnamed global package are visible to all user-defined namespaces without the need for an explicit **import** statement. To explicitly refer to a type declared in the unnamed global package, prefix the type name with "**::**".

import statements are not transitive. If package `B` imports package `A`, package `B` does not have unqualified access to contents declared in packages that `A` may have imported. Package `B` shall import those packages directly in order to have unqualified access to contents declared within them.

18.1.3.1 Syntax

```
import_stmt ::= import package_import_pattern ;
package_import_pattern ::= type_identifier [ package_import_qualifier ]
package_import_qualifier ::= package_import_wildcard | package_import_alias
package_import_wildcard ::= :: *
package_import_alias ::= as package_identifier
```

Syntax 88—Import statement

Note: *Package aliases* are described in 18.1.4.

Importing content from a **package** namespace using a wildcard only imports content from that exact namespace, and does not import content from *nested* namespaces.

Note that using a wildcard import on an outer package namespace, as shown with `p1::*` in the example below, allows inner package namespaces to be located without specifying the fully-qualified name of the namespace. In this example, **struct** `p1::p2::u` can be referenced as `p2::u` because the elements of `p1` are imported with a wildcard import.

Copyright © 2026 Accellera. All rights reserved.

```
package p1 {
    struct s { }
    package p2 {
        struct u { }
    }
}

struct t { }
struct s { }

package top {
    import p1::*;
    struct my_s {
        s     v1; // Resolves to p1::s
        ::s   v2; // Explicit reference to ::s
        t     v3; // Resolves to ::t
        p2::u v4; // Resolves to p1::p2::u
    }
}
```

Example 258—Importing the name of a nested package

18.1.4 Package aliases

The use of nested namespaces benefits from the ability to define a named *alias* for a given namespace. This is used when it is necessary to disambiguate between content declared in different namespaces and it is undesirable to use the fully-qualified name of the namespace. The syntax for declaring a package alias is shown in Syntax 88.

A namespace alias is only visible in the lexical scope (e.g., a package declaration statement) in which it appears. It is a name resolution shortcut, and does not introduce a new entity into the scope in which it is specified.

In the example below, this means that `p1` and `p2` are not visible in the scope of any other declaration statement of `consumer_pkg`. `p1` and `p2` may not be referenced from outside the package (e.g., as `consumer_pkg::p1`). Wildcard-importing `consumer_pkg` into another package namespace does not make symbols `p1` and `p2` visible in that namespace.

Copyright © 2026 Accellera. All rights reserved.

```
package pkg1::a::b::c {
    struct my_s {}
}

package pkg2::d::e::f {
    struct my_s {}
}

package consumer_pkg {
    import pkg1::a::b::c as p1;
    import pkg2::d::e::f as p2;
    struct s {
        p1::my_s            v1_1; // Refers to pkg1::a::b::c::my_s
        pkg1::a::b::c::my_s v1_2; // v1_1 and v1_2 have the same type
        p2::my_s            v2;   // Refers to pkg2::d::e::f::my_s
    }
}
```

Example 259—Package alias

A package alias shall not have the same name as a package name added to the same namespace in previous or current source units. However, it shall be legal to add a package name with the same name as the package alias in subsequent source units. In addition, two package aliases defined in the same lexical scope shall not have the same name.

```
package P {
    package foo {}
    package bar {}
}

package P {
    import bar as foo;       // Error: P already has a package named 'foo'
    import foo as my_alias;
    import bar as my_alias;  // Error: cannot define two aliases named
                             // 'my_alias' in the same scope
}
```

Example 260—Illegal package alias declarations

18.2 Declaration and reference ordering

Elements may be referenced after their declaration, within the same source unit or in a subsequent source unit. PSS also enables referencing most elements prior to their declaration within the same source unit, but places stronger ordering requirements on some elements. The following apply:

a) A variable declared and referenced within a procedural block may only be referenced after its declaration.

b) An action handle declared and referenced within an **activity** block may only be referenced after its declaration.

c) A constant or enum item may be referenced in the initialization assignment expression of another constant only after its declaration.

d) A constant declared within a type may reference type-level and package-level constants in its initialization assignment expression. A package-level constant may only reference other package-level constants in its initialization assignment expression.

Copyright © 2026 Accellera. All rights reserved.

18.2.1 Examples

In the example below, `file1.pss` (the first source unit) declares a **component** named `lib_base_c`. `file2.pss` (the second source unit) declares a type `my_base_c` that inherits from `lib_base_c`, so `file1.pss` shall be processed before `file2.pss`. However, within `file2.pss`, the declaration of `my_a_c` that refers to `my_base_c` as a supertype may be placed either before or after the declaration of `my_base_c`.

```
// Source Unit 1 (file1.pss)
component lib_base_c { /* ... */ }

// Source Unit 2 (file2.pss)
component my_a_c : my_base_c { /* ... */ }

component my_base_c : lib_base_c { /* ... */ }
```

Example 261—Reference to a previous source unit

In the example below, **action** `pss_top::entry` declares a field named `val` that is referenced in the **constraint** `val_c`. Field `val` may be declared before or after the **constraint** that references it.

```
component pss_top {
    action entry {
        constraint val_c {
            val < 10;
        }

        rand bit[4]    val;
    }
}
```

Example 262—Reference to a later-declared action field

In the example below, a local variable is declared within an **exec** block. As per requirement a) above, the variable `val` may only be referenced after it is declared.

```
function int get_val();

component pss_top {
    exec init_up {
        int val;
        val = get_val();
    }
}
```

Example 263—Reference to local variable after declaration

In the example below, constants are declared and referenced in initialization expressions of other constants. As per requirement c) above, a constant shall be declared prior to its reference in an initialization expression of a constant or in a type-width expression. Consequently, it is an error to reference the yet undeclared constant `C` in the initialization expression for `A`. It is legal to reference the previously declared constant `A` in the initialization expression for `B`.

Copyright © 2026 Accellera. All rights reserved.

```
package my {
  const int A = C /* Error: C is not yet declared */;
  const int B = A + 2;
  const int C = 3;
}
```

Example 264—Initialization of constants

18.3 Name resolution

For the purpose of the following description, the term *namespace* refers to either a **package** or a type (e.g., **component**, **struct**) under which static members (types, static constants, static functions, and enum items) may be declared.

The members of a package namespace include the members declared in the union of all the package definition statements of that package (see 18.1.1.1). The visible members of a type namespace include the members declared in the union of the type's initial definition and all visible extensions of the type (see 17.2),

Members of PSS namespaces shall have unique names in the context of their namespace, but members may have the same name if declared under different namespaces.

Types can be referenced in different contexts, such as declaring a variable, extending a type, or inheriting from a type. In all cases, a qualified name of the type can be used, using the scope operator **::**.

Constants, static functions, and enum items can be referenced in expression contexts. In these cases too, a qualified name can be used, using the scope operator.

Informally, unqualified entity names can be used in the following cases:

- when referencing an entity that was declared in the same namespace or in an enclosing namespace.
- when referencing an entity that was declared in a **package** imported into a logical scope enclosing the reference.

Precedence is given to the current namespace scope; explicit qualification can be used to override the precedence.

Formally, unqualified names are resolved using the following process, starting with step a, continuing with step b, and then step c, in the absence of resolution in previous steps:

a) If the reference occurs within an expression whose *expected type* is an enumeration type (see 8.4.3 for definition of expected type):
   1) Search enum items declared in the expected type's initial definition.
   2) Search enum items declared in the expected type's extensions that are defined under the current package or one of its containing packages (see 17.2), or in the expected type's extensions that are within a package wildcard-imported into a lexical scope enclosing the reference.

b) If the reference occurs within the definition of a type:
   1) Search members of the type declared in its initial definition.
   2) Search members of the type declared in its extensions that are defined under the current package or one of its containing packages (see 17.2), or in its extensions that are within a package wildcard-imported into a lexical scope enclosing the reference.
   3) If the type inherits from a supertype, search members declared in the supertype using the process described in steps 1 and 2. Repeat for all supertypes in the inheritance hierarchy.

Copyright © 2026 Accellera. All rights reserved.

4) If the scope is a component initial definition or extension:
   i) Search package members explicitly imported into the lexical scope of the initial definition or extension, respectively.
   ii) Search members of packages wildcard-imported into the lexical scope of the initial definition or extension, respectively.
5) If the type is an inner type (e.g., an **action** declared inside a **component**), search members declared in the outer type using the process described in steps 1 through 4 above.

c) Search **package** namespaces, starting with the package namespace of the immediate lexical scope and working outward along the package hierarchy. At each level, do the following:
   1) Search package members declared under all *package_declaration*s of the same package.
   2) If the reference is enclosed in a lexical package scope corresponding to the namespace being searched:
      i) If the package member being searched for is itself a package, search for a package alias name defined in the lexical scope of the corresponding *package_declaration* statement.
      ii) Search package members explicitly imported into the lexical scope.of the corresponding *package_declaration* statement.
      iii) Search members of packages wildcard-imported into the lexical scope.of the corresponding *package_declaration* statement.

A qualified name is composed of double-colon-separated elements. Qualified name elements are resolved by first applying the same process for unqualified names described above on the first element of the static path. Having resolved the first element to a certain **package**/type, the rest of the static path is used to access down from it.

18.3.1 Name resolution examples

In Example 265, `s` is declared in three places: imported package `P1`, encapsulating package `P2`, and nested component `C1`. The `s` referenced in nested component `C1` is resolved to the `s` locally defined in nested component `C1`. Using qualifiers, `P1::s` would be used to resolve to `s` in imported package `P1`, and `P2::s` would be used to resolve to `s` in encapsulating package `P2`.

```
package P1 {
  struct s {};
};

package P2 {
  struct s {};

  component C1 {
    import P1::*;
    struct s {};
    s f;
  };
};
```

Example 265—Name resolution to declaration in nested namespace

Copyright © 2026 Accellera. All rights reserved.

In Example 266, `s` is declared in two places: imported package `P1` and encapsulating package `P2`. The `s` referenced in nested component `C1` is resolved to the `s` defined in imported package `P1`. Using qualifiers, `P2::s` would be used to resolve to `s` in encapsulating package `P2`.

```
package P1 {
  struct s {};
};

package P2 {
  struct s {};

  component C1 {
    import P1::*;
    s f;
  };
};
```

Example 266—Name resolution to declaration in imported package in nested namespace

In Example 267, `s` is declared in two places: imported package `P1` and encapsulating package `P2`. The `s` referenced in nested component `C1` is resolved to the `s` defined in encapsulating package `P2`. Using qualifiers, `P1::s` would be used to resolve to `s` in package `P1` imported in encapsulating package `P2`.

```
package P1 {
  struct s {};
};

package P2 {
  import P1::*;
  struct s {};

  component C1 {
    s f;
  };
};
```

Example 267—Name resolution to declaration in encapsulating package

In Example 268, `s` is declared in one place: imported package `P1`. The `s` referenced in nested component `C1` is resolved to the `s` defined in package `P1` imported inside encapsulating package `P2`.

```
package P1 {
  struct s {};
};

package P2 {
  import P1::*;

  component C1 {
    s f;
  };
}
```

Example 268—Name resolution to declaration in imported package in encapsulating package

Copyright © 2026 Accellera. All rights reserved.

Example 269 shows a case where importing the encapsulating package has no effect on the resolution rules. `s` will resolve to the same `s` in `P2`.

```
package P1 {
  struct s {};
};

package P2 {
  import P1::*;
  struct s {};

  component C1 {
    import P2::*;
    s f;
  };
}
```

Example 269—Package import has no effect on name resolution

Example 270 shows a case where importing the encapsulating package does have effect on the resolution rules. `s` will resolve to `s` in `P1` due to the wildcard import of `P1`.

```
package P1 {
  struct s {}

  package P2 {
    struct s {}
    component C1 {
      import P1::*;
      s f;         // P1::s
      P2::s g;     // P1::P2::s
    }
  }
}
```

Example 270—Package import affects name resolution

In Example 271 below, `a_pkg` declares a **struct** `S1`, `b_pkg` imports content from `a_pkg`, and `b_pkg` declares a **struct** `S2` that inherits from `S1`. `pss_top` imports content from `b_pkg`.

- Line (1): `S2` is resolved via the import of `b_pkg`.
- Line (2): Imports are not transitive. Therefore, the import of `b_pkg` does not make content from `a_pkg` visible in **component** `pss_top`.
- Line (3): `S1` can be referenced with a fully-qualified type name, `a_pkg::S1`.
- Line (4): Importing a package does not introduce symbols into the importing namespace.

Copyright © 2026 Accellera. All rights reserved.

```
package a_pkg {
  struct S1 { }
}

package b_pkg {
  import a_pkg::*;
  struct S2 : S1 { }
}

component pss_top {
  import b_pkg::*;
  S2         s2_i0; // (1) OK
  S1         s1_i1; // (2) Error: S1 is not made visible
                    //            by importing b_pkg
  a_pkg::S1  s1_i2; // (3) OK: S1 is declared in a_pkg
  b_pkg::S1  s1_i3; // (4) Error: import of a_pkg in b_pkg
                    //            does not make S1 a b_pkg member
};
```

Example 271—Package import is not a declaration

Example 272 demonstrates the use of qualified and unqualified enum item references. The unqualified references are resolved based on the expected type in context, namely the type of the expression on the other side of the equality operator and on the left-hand side of the **in** operator.

```
component my_ip_c {
  enum mode_e {A, B, C, D};
  action my_op {
    rand mode_e mode;
  }
}

component pss_top {
  my_ip_c my_ip;
  action test {
    my_ip_c::my_op op;
    constraint op.mode == my_ip_c::mode_e::A;
    constraint op.mode == A;
    constraint op.mode in [A, C, D];

    activity {
      op;
    }
  }
}
```

Example 272—Resolution of enum item references

Example 273 demonstrates how name resolution is affected by using package aliases. `P2::s` is resolved to `P3::P4::s` and not to `P1::P2::s`, because the package alias takes precedence over the wildcard import in resolving `P2`.

Copyright © 2026 Accellera. All rights reserved.

```
package P1 {
  package P2 {
    struct s {}
  }
}

package P3 {
  package P4 {
    struct s {}
  }
}

component pss_top {
  import P1::*;
  import P3::P4 as P2;
  action test {
    P2::s f;
  }
}
```

Example 273—Resolution in presence of package alias

Copyright © 2026 Accellera. All rights reserved.

19. Conditional code processing

It is often useful to conditionally process portions of a PSS model based on some configuration parameters. This clause details a **compile if** construct that can be evaluated as part of the elaboration process.

19.1 Overview

This section covers general considerations for using compile statements.

19.1.1 Statically-evaluated statements

A *statically-evaluated statement* marks content that may or may not be elaborated. The description within a statically-evaluated statement shall be syntactically correct, but need not be semantically correct when the static scope is disabled for evaluation.

A statically-evaluated statement may specify a block of statements. However, this does not introduce a new scope in the resulting description.

19.1.2 Elaboration procedure

Compile statements are processed top-to-bottom within a given source unit. The following steps are performed in processing source code in the presence of conditional compilation directives:

a) Syntactic code analysis is performed.
b) Compile-time expressions are evaluated in order within the following contexts:
    1) **static const** initializers
    2) **compile if** conditions (see 19.2)

    These expressions are evaluated based on types and static constants declared:
    1) Unconditionally, or in an enabled **compile if** branch, within a previously-processed source unit
    2) Unconditionally, or in an enabled **compile if** branch, previously processed within the current source unit
c) Globally-visible content and the content within enabled **compile if** branches is elaborated.

19.1.3 Compile-time expressions

The value of any **compile if** expressions shall be determinable at compile time. Because **compile if** statements are evaluated early in PSS source processing, only types and constants declared in **package** scopes may be referenced. Types and constants declared in type scopes (e.g., an **action** type declared within a **component** type) may not be referenced.

The example below highlights the reference rules for conditional compilation directives:

a) Conditional compilation directives are evaluated based on previously defined elements.
    1) Consequently, the first directive (`compile has(s)`) evaluates true because `p1::s` is visible at this point in the evaluation.
    2) The second directive (`compile has(t)`) also evaluates true because `p2::t` has been previously declared in the source unit.
b) Conditional compilation directives may not reference inner members of types. Consequently, attempting to reference `t::A` is an error, since `t` is a type and `A` is an inner member of type `t`.

Copyright © 2026 Accellera. All rights reserved.

```
package p1 {
   struct s {
      static const int A = 3;
   };
};

package p2 {
   import p1::*;

   // derived from p2::s defined later in this file
   struct t : s { };

   // evaluates to true because such a type has been previously defined,
   // namely p1::s
   compile if (compile has (s)) { … }

   // evaluates to true because such a type has been previously defined,
   // namely p2::t (even though its supertype is not yet known)
   compile if (compile has (t)) { … }

   // Illegal! Cannot reference a member of a struct in compile-if context
   compile if (t::A == 2) { … }

   struct s {};
}
```

Example 274—Conditional compilation evaluation

19.2 Compile if

19.2.1 Scope

compile if statements may appear in the following scopes:

- Global/package
- Action
- Component
- Struct
- Procedural Scopes (Execs[1] and Functions)
- Constraints
- Covergroups
- Overrides

[1]Excluding target-template exec-body blocks

Copyright © 2026 Accellera. All rights reserved.

19.2.2 Syntax

Syntax 89 shows the grammar for a **compile if** statement.

```
package_body_compile_if ::= compile if ( constant_expression )
    package_body_compile_if_item [ else package_body_compile_if_item ]
monitor_body_compile_if ::= compile if ( constant_expression )
    monitor_body_compile_if_item [ else monitor_body_compile_if_item ]
action_body_compile_if ::= compile if ( constant_expression )
    action_body_compile_if_item [ else action_body_compile_if_item ]
component_body_compile_if ::= compile if ( constant_expression )
    component_body_compile_if_item [ else component_body_compile_if_item ]
struct_body_compile_if ::= compile if ( constant_expression )
    struct_body_compile_if_item [ else struct_body_compile_if_item ]
procedural_compile_if ::= compile if ( constant_expression )
    procedural_compile_if_stmt [ else procedural_compile_if_stmt ]
constraint_body_compile_if ::= compile if ( constant_expression )
    constraint_body_compile_if_item [ else constraint_body_compile_if_item ]
covergroup_body_compile_if ::= compile if ( constant_expression )
    covergroup_body_compile_if_item [ else covergroup_body_compile_if_item ]
override_compile_if ::= compile if ( constant_expression )
    override_compile_if_stmt [ else override_compile_if_stmt ]
package_body_compile_if_item ::= { { package_body_item } }
action_body_compile_if_item ::= { { action_body_item } }
component_body_compile_if_item ::= { { component_body_item } }
struct_body_compile_if_item ::= { { struct_body_item } }
procedural_compile_if_stmt ::= { { procedural_stmt } }
constraint_body_compile_if_item ::= { { constraint_body_item } }
covergroup_body_compile_if_item ::= { { covergroup_body_item } }
override_compile_if_stmt ::= { { override_stmt } }
```

Syntax 89—Compile if declaration

19.2.3 Examples

Example 275 shows an example of conditional processing if PSS were to use C pre-processor directives. If the `PROTOCOL_VER_1_2` directive is defined, then action `new_flow` is evaluated. Otherwise, action `old_flow` is processed.

NOTE—Example 275 is only shown here to illustrate the functionality of C pre-processor directives in a familiar format. It is not part of PSS.

Copyright © 2026 Accellera. All rights reserved.

```
#ifdef PROTOCOL_VER_1_2
action new_flow {
  activity { ... }
}
#else
action old_flow {
  activity { ... }
}
#endif
```

Example 275—Conditional processing (C pre-processor)

Example 276 shows a PSS version of Example 275 using a **compile if** statement instead.

```
package config_pkg {
  const bool PROTOCOL_VER_1_2 = false;
}
compile if (config_pkg::PROTOCOL_VER_1_2) {
  action new_flow {
    activity { ... }
  }
} else {
  action old_flow {
    activity { ... }
  }
}
```

Example 276—Conditional processing (compile if)

When the *true* case is triggered, the code in Example 276 is equivalent to:

```
action new_flow {
    activity { ... }
  }
```

When the *false* case is triggered, the code in Example 276 is equivalent to:

```
action old_flow {
    activity { ... }
  }
```

19.3 Compile has

compile has allows conditional elaboration to reason about the existence of types and constants. The **compile has** expression evaluates to *true* if a type or constant has been previously declared unconditionally or within an enabled conditional block (see 19.1.2); otherwise, it evaluates to *false*.

19.3.1 Syntax

Syntax 90 shows the grammar for a **compile has** expression.

Copyright © 2026 Accellera. All rights reserved.

```
compile_has_expr ::= compile has ( static_ref_path )
static_ref_path ::= [ :: ] { type_identifier_elem :: } member_path_elem
```

Syntax 90—Compile has expression

19.3.2 Examples

Example 277 checks whether the `config_pkg::PROTOCOL_VER_1_2` field exists and tests its value if it does. In this example, `old_flow` will be used because `config_pkg::PROTOCOL_VER_1_2` does not exist.

```
package config_pkg {
}

compile if (compile has(config_pkg::PROTOCOL_VER_1_2) &&
            config_pkg::PROTOCOL_VER_1_2) {
  action new_flow {
    activity { ... }
  }
} else {
  action old_flow {
    activity { ... }
  }
}
```

Example 277—Compile has

Example 278 is composed of a single source unit.

— The first top-level **compile if** block checks for the existence of `X`. This evaluates to *false*, since `X` is only subsequently declared within the source unit.

— The second top-level **compile if** block checks for the non-existence of `Y`. This evaluates to *true*, since `Y` was not previously declared (the first **compile if** block was not expanded). As a consequence, `Y` is declared with a value of `0`.

```
compile if (compile has(X)) {
    const int Y = 2;
    compile if (compile has(Y)) {
        const int Z;
    }
}

const int X = 1;

compile if (!(compile has(Y))) {
    const int Y=0;
} else {
    compile if (compile has(Z)) {
        const int A;
    }
}
```

Example 278—Nested conditions

Copyright © 2026 Accellera. All rights reserved.

19.4 Compile assert

compile assert assists in flagging errors when the source is incorrectly configured. This construct is evaluated during elaboration. A tool shall report a failure if *constant_expression* does not evaluate to *true*, and report the user-provided message, if specified.

19.4.1 Syntax

Syntax 91 shows the grammar for a **compile assert** statement.

```
compile_assert_stmt ::= compile assert ( constant_expression [ , string_literal ] ) ;
```

Syntax 91—Compile assert statement

19.4.2 Example

Example 279 shows a **compile assert** example.

```
compile if (compile has(FIELD2)) {
  static const FIELD1 = 1;
}

compile if (compile has(FIELD1)) {
  static const FIELD2 = 2;
}
compile assert(compile has(FIELD1), "FIELD1 not found");
```

Example 279—Compile assert

Copyright © 2026 Accellera. All rights reserved.

20. Test realization

A PSS model interacts with foreign languages in order to drive, or bring about, the behaviors that leaf-level actions represent in a test scenario. This is done by calling application programming interfaces (APIs) available in the execution environment, or generating foreign language code that executes as part of the test. In addition, external code, such as reference models and checkers, may be used to help compute stimulus values or expected results during stimulus generation.

The constructs defined in this clause specify the fundamental semantics of test realization. The Core Library, described in Clause 21, builds on these foundational semantics and provides standardized constructs that make use of the realization model described herein. Accordingly, this chapter may reference Core Library elements where their definitions rely on these semantics. The normative definitions of those library constructs are provided in the Core Library clause.

The platform on which test generation takes place is generally referred to as the *solve platform*, while the platform on which test execution takes place is called the *target platform*.

Logic used to help compute stimulus values is coded using *procedural constructs* (see 20.7), possibly invoking a foreign procedural interface on the solve platform (see 20.4). The implementation of runtime behavior of leaf-level actions can similarly be specified with procedural constructs, possibly invoking a foreign procedural interface on the target platform or invoking *target template functions* (see 20.6). Alternatively, implementation of actions and other scenario entities can be specified as *target code template blocks* (see 20.5). In all cases, the constructs for specifying implementation of PSS entities are called *exec block*s.

Functions can be defined in PSS as a means to factor out and reuse portable procedural logic required for the implementation of scenario entities in *exec block*s (see 20.3). Functions may take parameters and optionally return a result value. Like *exec block*s, functions are defined in terms of procedural constructs or as target code templates.

20.1 Exec blocks

*exec block*s provide a mechanism for associating specific functionality with a **component**, an **action**, a flow/resource object, or a **struct** (see Syntax 92). A number of *exec block* kinds are used to implement scenario entities.

- **init_down** and **init_up** *exec block*s allow component data fields to be assigned a value as the component tree is being elaborated (see 9.1.4).
- **body** *exec block*s specify the actual runtime implementation of atomic actions.
- **pre_solve** and **post_solve** *exec block*s of **action**s, flow/resource objects, and **struct**s are a way to involve arbitrary computation as part of the scenario solving.
- Other **exec** kinds serve more specific purposes in the context of pre-generated test code and auxiliary files.

Copyright © 2026 Accellera. All rights reserved.

20.1.1 Syntax

```
exec_block_stmt ::=
    exec_block
  | target_code_exec_block
  | target_file_exec_block
exec_block ::= exec exec_kind { { procedural_stmt } }
exec_kind ::=
    pre_solve
  | post_solve
  | pre_body
  | body
  | header
  | declaration
  | run_start
  | run_end
  | init_down
  | init_up
exec_super_stmt ::= super ;
target_code_exec_block ::= exec exec_kind language_identifier =
    [exec_block_tag :] string_literal ;
target_file_exec_block ::= exec file filename_string = [exec_block_tag :] string_literal ;
    filename_string ::= string_literal
```

Syntax 92—Exec block declaration

The following also apply:

a) *exec block* content is given in one of two forms: as a sequence of procedural constructs (possibly involving foreign function calls) or as a text segment of target code parameterized with PSS attributes.

b) In either case, a single *exec block* is always mapped to implementation in no more than one foreign language.

c) In the case of a target-template block, the target language shall be explicitly declared; however, when using procedural constructs, the corresponding language may vary.

d) Multiple *exec block*s of the same kind may be declared in a given definition scope. If multiple *exec block*s of the same kind are declared in a given definition scope, they shall be considered as a single *exec block* of the given kind, processed in source order.

e) The execution of solve-time exec blocks shall be atomic with respect to shared component data accesses.

20.1.2 Exec block kinds

The following list describes the different *exec block* kinds:

— **init_down/init_up**—valid in **component** types. The **init_down** and **init_up** blocks are used to assign values to component attributes and to initialize foreign language objects. Component

Copyright © 2026 Accellera. All rights reserved.

init_down and **init_up** blocks are called before the scenario root action's **pre_solve** block is invoked. **init_down** and **init_up** blocks may not call target template functions.

1) **init_down**—Starting with the root component, **init_down** blocks are evaluated top-down for each component in the hierarchy. The relative order of evaluating **init_down** blocks for components at the same level of hierarchy is undefined. For any component, the **init_down** block shall be evaluated before its **init_up** block is evaluated.
2) **init_up**—For a leaf-level component (i.e., one that does not instantiate any subcomponents), the **init_up** block shall be evaluated after its **init_down** block (if any). A parent component's **init_up** block shall be evaluated only after all subcomponent **init_up** blocks have been evaluated.

— **pre_solve**—valid in **action**, flow/resource object, and **struct** types. The **pre_solve** block is processed prior to solving of random-variable relationships in the PSS model. **pre_solve** *exec block*s are used to initialize non-random variables that the solve process uses. See also 13.4.12.

— **post_solve**—valid in **action**, flow/resource object, and **struct** types. The **post_solve** block is processed after random-variable relationships have been solved. The **post_solve** *exec block* is used to compute values of non-random fields based on the solved values of random fields. See also 13.4.12.

— **pre_body**—valid in **action**, flow/resource object, and **struct** types. The **pre_body** block is an *exec block* evaluated on the solve platform that is evaluated after **exec post_solve** and before **exec body** is evaluated as part of the test realization process. It is evaluated after executor assignments and memory allocations are completed for the given action, but before code is generated to represent the body block. Solve functions may be called in this *exec block*, as well as the **`executor()`**, **`addr_value_solve()`**, and **`addr_value_abs()`** functions.

— **body**—valid in **action** types. The **body** block constitutes the implementation of an atomic **action**. The **body** block of each **action** is invoked in its respective order during the execution of a scenario—after the **body** blocks of all predecessor **actions** complete. Execution of an **action**'s **body** may be logically time-consuming and concurrent with that of other actions. In particular, the invocation of **exec** blocks of **action**s with the same set of scheduling dependencies logically takes place at the same time. Implementation of the standard should guarantee that executions of **exec** blocks of same-time **action**s take place as close as possible.

— **header**—valid in **component**, **action**, flow/resource object, and **struct** types. The **header** block specifies top-level statements for header declarations presupposed by subsequent code blocks of the context action or object. Examples are **`'#include'`** directives in C, or forward function or class declarations. For more information on evaluation order of target execs see 20.1.5.

— **declaration**—valid in **component**, **action**, flow/resource object, and **struct** types. The **declaration** block specifies declarative statements used to define entities that are used by subsequent code blocks. Examples are the definition of global variables or functions. For more information on evaluation order of target execs see 20.1.5.

— **run_start**—valid in **component**, **action**, flow/resource object, and **struct** types. The **run_start** block is a procedural code block to be executed before any **body** block of the scenario is invoked on the same executor. It is used typically for one-time test bring-up and configuration required by the context action or object. **exec run_start** is restricted to the pre-generation flow (see Table 29). For more information on evaluation order of target execs see 20.1.5.

— **run_end**—valid in **component**, **action**, flow/resource object, and **struct** types. The **run_end** block is a procedural code block to be executed after all **body** blocks of the scenario are completed on the same executor. It is used typically for test bring-down and post-run checks associated with the context action or object. **exec run_end** is restricted to pre-generation flow (see Table 29). For more information on evaluation order of target execs see 20.1.5.

exec header and **declaration** blocks shall only be specified in terms of target code templates. All other target **exec** kinds may be specified in terms of procedural constructs or target code templates.

Copyright © 2026 Accellera. All rights reserved.

20.1.3 Examples

In Example 280, the **init_up** *exec block*s are evaluated in the following order:

a) `init_up in pss_top.s1`
b) `init_up in pss_top.s2`
c) `init_up in pss_top`

This results in the **component** fields having the following values:

a) `s1.base_addr=0x2000 (init_up in pss_top overwrote the value set by init_up in sub_c)`
b) `s2.base_addr=0x1000 (value set by init_up in sub_c)`

```
component sub_c {
  int base_addr;

  exec init_up {
    base_addr = 0x1000;
  }
};

component pss_top {
  sub_c s1, s2;

  exec init_up {
    s1.base_addr = 0x2000;
  }
};
```

Example 280—Data initialization in a component

In Example 281, the **init_down** and **init_up** blocks will be evaluated in the following order:

- `init_down in T`
- `init_down in T.c1`
- `init_down in T.c2`
- `init_up in T.c1`
- `init_up in T.c2`
- `init_up in T`

Copyright © 2026 Accellera. All rights reserved.

```
component C {
    exec init_down {
    }
    exec init_up {
    }
}

component T {
    C  c1, c2;
    exec init_down {
    }
    exec init_up {
    }
}
```

Example 281—Init_down and init_up exec blocks

A diagram of the example is shown below:

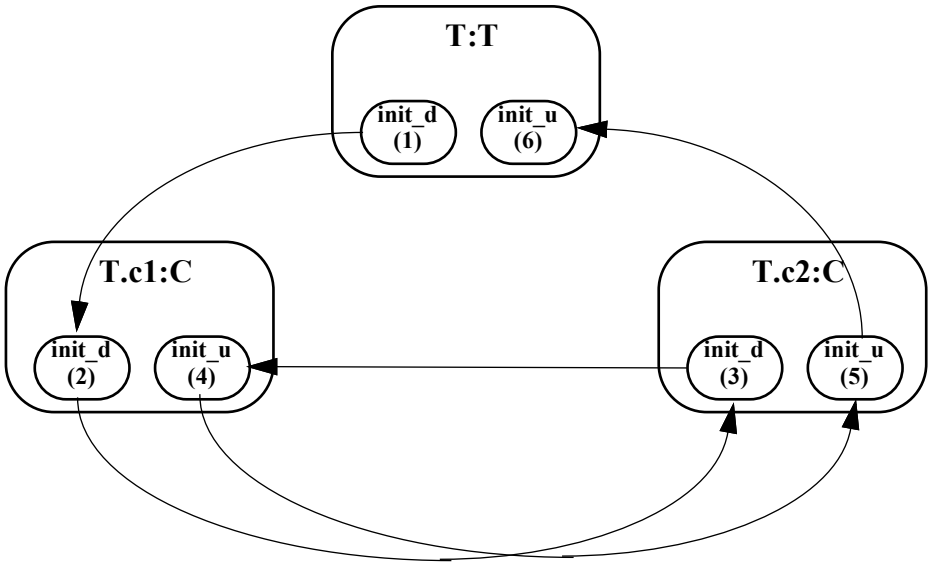


Figure 51—Order of invocation of init_down and init_up exec blocks

The order of initialization calls is annotated on each of the **init_d(own)** and **init_u(p)** blocks. Note that **init_down** in `T` is called first, followed by **init_down** in `T.c1`, etc.

Note that a tool is free to execute the **exec init_down** and **init_up** blocks of sibling instances in arbitrary order. For example, while the diagram above shows **init_down** in `T.c1` executing before **init_down** in `T.c2`, the opposite order is also correct. The key requirements are that the **exec init_down** block of a parent component instance (e.g., `T`) execute before the **exec init_down** block of any child **component** instances, and that the **exec init_up** block of a parent component instance (e.g., `T`) execute after all **exec init_up** blocks of child component instances have executed. This implies that the following ordering of execution is also legal:

- `init_down in T`
- `init_down in T.c1`

Copyright © 2026 Accellera. All rights reserved.

- `init_up in T.c1`
- `init_down in T.c2`
- `init_up in T.c2`
- `init_up in T`

In Example 282, component `pss_top` contains two instances of component `sub_c`, named `s1` and `s2`. Component `sub_c` contains a data field named `base_addr` that controls the value to function `activate()` when action `A` is traversed.

During construction of the component tree, component `pss_top` sets `s1.base_addr=0x1000` and `s2.base_addr=0x2000`.

Action `pss_top::entry` traverses action `sub_c::A` twice. Depending on which component instance `sub_c::A` is associated with during traversal, it will cause `sub_c::A` to be associated with a different `base_addr`.

- If `sub_c::A` executes in the context of `pss_top.s1`, `sub_c::A` uses `0x1000`.
- If `sub_c::A` executes in the context of `pss_top.s2`, `sub_c::A` uses `0x2000`.

```
component sub_c {
    bit[32] base_addr = 0x1000;
    action A {
        exec body {
            // reference base_addr in context component
            activate(comp.base_addr + 0x10);
                            // activate() is an imported function
        }
    }
}

component pss_top {
    sub_c s1, s2;
    exec init_up {
        s1.base_addr = 0x1000;
        s2.base_addr = 0x2000;
    }
    action entry {
        sub_c::A a;
        activity {
            repeat (2) {
                a;  // Runs sub_c::A with 0x1000 as base_addr when
                    // associated with s1
                    // Runs sub_c::A with 0x2000 as base_addr when
                    // associated with s2
            }
        }
    }
}
```

Example 282—Accessing component data field from an action

For additional examples of *exec block* usage, see 20.2.7.

Copyright © 2026 Accellera. All rights reserved.

20.1.4 Exec block evaluation with inheritance and extension

Both inheritance and type extension can impact the behavior of *exec block*s. See also 17.1 and 17.2.

20.1.4.1 Inheritance and shadowing

*exec block*s are considered to be *virtual*, in that a derived type that defines an *exec block* completely replaces the behavior of any same-kind *exec block* (e.g., **body**) specified by its base type. Procedural *exec block*s may include the "**`super;`**" statement, which will execute the contents of the corresponding base-type *exec block*(s) at that point (see 20.1.4.2).

The following examples use the core library target function **`message()`** to add a formatted line as a message to the execution log, at **`LOW`** verbosity. In Example 283, action `B` inherits from action `A` and shadows the **pre_solve** and **body** *exec block*s defined by action `A`.

```
action A {
  int a;

  exec pre_solve {
    a=1;
  }
  exec body {
    std_pkg::message(LOW,"Hello from A %d", a);
  }
}

action B : A {
  exec pre_solve {
    a=2;
  }
  exec body {
    std_pkg::message(LOW,"Hello from B %d", a);
  }
}
```

Example 283—Inheritance and shadowing

When an instance of action `B` is evaluated, the following is printed:

```
Hello from B 2
```

20.1.4.2 Using super

Specifying "**`super;`**" as a statement in a subtype executes the behavior of the same-kind procedural *exec block*(s) from the base type. The "**`super;`**" statement shall not be allowed in a target-template *exec* block. The "**`super;`**" statement shall not be allowed in function bodies.

In Example 284, both `A1` and `A2` inherit from action `A`. Both execute the **pre_solve** *exec block* inherited from `A`. `A1` invokes the **body** behavior of `A`, then displays an additional statement. `A2` displays an additional statement, then invokes the **body** behavior of `A`.

Copyright © 2026 Accellera. All rights reserved.

```
action A {
  int a;

  exec pre_solve {
    a=1;
  }
  exec body {
    std_pkg::message(LOW,"Hello from A %d", a);
  }
}

action A1 : A {
  exec body {
    super;
    std_pkg::message(LOW,"Hello from A1 %d", a);
  }
}
action A2 : A {
  exec body {
    std_pkg::message(LOW,"Hello from A2 %d", a);
    super;
  }
}
```

Example 284—Using super

When an instance of `A1` is evaluated, the following is printed:

```
Hello from A 1
Hello from A1 1
```

When an instance of `A2` is evaluated, the following is printed:

```
Hello from A2 1
Hello from A 1
```

20.1.4.3 Type extension

Type extension enables additional features to be contributed to **action**, **component**, and **struct** types. Type extension is additive and all *exec block*s contributed via type extension are evaluated, along with *exec block*s specified within the initial definition. First, the initial definition's *exec block*s (if any) are evaluated. Next, the *exec block*s (if any) contributed via type extension are evaluated, in the order that they are processed by the PSS processing tool.

Copyright © 2026 Accellera. All rights reserved.

In Example 285, a type extension contributes an *exec block* to action `A1`.

```
action A {
  int a;

  exec pre_solve {
    a=1;
  }
  exec body {
    std_pkg::message(LOW,"Hello from A %d", a);
  }
}

action A1 : A {
  exec body {
    super;
    std_pkg::message(LOW,"Hello from A1 %d", a);
  }
}

extend action A1 {
  exec body {
   message(LOW,"Hello from A1 extension %d", a);
  }
}
```

Example 285—Type extension contributes an exec block

When an instance of `A1` is evaluated, the following is printed:

```
Hello from A 1
Hello from A1 1
Hello from A1 extension 1
```

Copyright © 2026 Accellera. All rights reserved.

In Example 286, two *exec block*s are added to action `A1` via extension.

```
action A {
  int a;

  exec pre_solve {
    a=1;
  }
  exec body {
    std_pkg::message(LOW,"Hello from A %d", a);
  }
}

action A1 : A {
  exec body {
    super;
    std_pkg::message(LOW,"Hello from A1 %d", a);
  }
}

extend action A1 {
  exec body {
    std_pkg::message(LOW,"Hello from A1(1) extension %d", a);
  }
}

extend action A1 {
  exec body {
    std_pkg::message(LOW,"Hello from A1(2) extension %d", a);
  }
}
```

Example 286—Exec blocks added via extension

If the PSS processing tool processes the first extension followed by the second extension, then the following is produced:

```
Hello from A 1
Hello from A1 1
Hello from A1(1) extension 1
Hello from A1(2) extension 1
```

If the PSS processing tool processes the second extension followed by the first extension, then the following is produced:

```
Hello from A 1
Hello from A1 1
Hello from A1(2) extension 1
Hello from A1(1) extension 1
```

20.1.5 Evaluation order of target exec blocks

Target exec kinds refer to exec *body*, *run_start*, *run_end*, *header*, and *declaration*.

During generation, exec blocks of actions participating in the scenario, exec blocks of their flow/resource objects, and exec blocks of active components are instantiated in the respective execution unit. A component

Copyright © 2026 Accellera. All rights reserved.

is considered active if it has an action that is part of the scenario. For more information on how the execution unit is set, see 21.8.

20.1.5.1 Evaluation of run_start and run_end

run_start and *run_end* blocks are instantiated and hence executed sequentially in the respective executors. The order of evaluation between exec kinds or between exec blocks targeted for different executors is not defined. For more information about assignment of executors see 21.7.2.

The exec blocks of the active components are executed first followed by the blocks of the actions and their flow/resource objects.

20.1.5.1.1 Components

run_start and *run_end* exec blocks are evaluated per executor following the component instance tree hierarchy, with special ordering semantics for execution_unit and executor. Order across executors is undefined. *run_start* and *run_end* exec blocks are applied strictly to component instances and not on component instance references.

For *run_start*, execution proceeds in the following order: first the execution_unit *run_start*, then the executor *run_start*, followed by component-level *run_start* exec blocks. Component-level *run_start* blocks execute top down, beginning at the root component and visiting active components in a depth-first pre-order traversal (parent before children).

For *run_end*, execution proceeds in the reverse order: component-level *run_end* exec blocks execute first, followed by executor *run_end*, and finally, the execution_unit *run_end* block. Component-level *run_end* executes bottom up, visiting active components in a depth-first post-order traversal (children before parent).

When *run_start* and *run_end* exec blocks are implemented in an execution unit, they are duplicated for each active executor associated with that execution unit and executed as explained above.

20.1.5.1.2 Actions

Exec blocks of actions are evaluated per executor instance according to the scheduling in the scenario.

For a given exec kind and execution unit, the order in which exec blocks of actions are evaluated conforms to their scheduling and flow order. An action that inputs a flow/resource object will always have its exec block of a specific kind evaluated after the action that outputs the flow/resource object and similarly for an action scheduled after another. For two actions scheduled in concurrent threads with no flow/resource object exchange between them, the exec evaluation order is undefined.

20.1.5.1.3 Buffer and state objects

The order in which *run_start* and *run_end* exec blocks of buffer and state objects are evaluated conforms to their scheduling order and the executor on which they are scheduled. When a producer action, a buffer or state object, and a consumer action are scheduled on the same executor, their exec *run_start* and exec *run_end* blocks are evaluated according to the scheduling order, with the flow object exec block processed between the producer and consumer actions exec blocks. No ordering relationship is implied when these constructs are associated with different executors.

20.1.5.1.4 Stream objects

No ordering relationship is implied for stream objects because of the concurrency nature of streams.

Copyright © 2026 Accellera. All rights reserved.

20.1.5.1.5 Resource objects

The *run_start* and *run_end* exec blocks of resource objects are evaluated before the respective exec block of the locking action.

20.1.5.1.6 Struct types

The *run_start* and *run_end* exec blocks of struct types are evaluated according to their declaration order in the containing entity.

20.1.5.2 Generation of header and declaration

The ordering for *header* and *declaration* exec blocks is similar to *run_start* exec.

When implementing *header* and *declaration* exec blocks in an execution unit, they are generated once per execution unit. Generation order is still according to the component hierarchy.

20.1.5.3 Examples

```
component A_executor_c : executor_c<> {
    exec run_start { ... }
}
component C_execution_unit : target_execution_unit_c {
    exec run_start { ... }
}
component my_comp {
    A_executor_c t;
    C_execution_unit e;

    exec run_start { ... }
    exec init_down {
        set_executor(t);
        t. set_target_execution_unit(e);
    }

    action act {
        exec run_start { ... }
        exec body { ... }
    }
}

component pss_top {
    my_comp c1,c2;
    action test {
        activity {
            parallel {
                do my_comp::act with {comp == pss_top.c1};
                do my_comp::act with {comp == pss_top.c2};
            }
        }
    }
}
```

Example 287—Order of evaluation for run_start exec

Copyright © 2026 Accellera. All rights reserved.

Example 287 demonstrates the order in which *run_start* exec will be evaluated depending on the component hierarchy:

```
pss_top
|── c1 : my_comp
|    |── t : A_executor_c
|    |── e : C_execution_unit
|    |── action act
|── c2 : my_comp
     |── t : A_executor_c
     |── e : C_execution_unit
     |── action act
```

Executor 1:

```
1. run_start of pss_top.c1.e
2. run_start of pss_top.c1.t
3. run_start of pss_top.c1
4. run_start of action act
```

Executor 2:

```
1. run_start of pss_top.c2.e
2. run_start of pss_top.c2.t
3. run_start of pss_top.c2
4. run_start of action act
```

Executor 1 and Executor 2 can be interleaved, for example:

```
1. Executor 1: run_start of pss_top.c1.e
2. Executor 1: run_start of pss_top.c1.t
3. Executor 2: run_start of pss_top.c2.e
4. Executor 1: run_start of pss_top.c1
5. Executor 2: run_start of pss_top.c2.t
6. Executor 2: run_start of pss_top.c2
7. Executor 2: run_start of action act
8. Executor 1: run_start of action act
```

The same output would be received if the execution unit is instantiated under pss_top:

```
pss_top
  |── c1 : my_comp
  |    |── t : A_executor_c
  |    |── action act
  |── c2 : my_comp
       |── t : A_executor_c
       |──action act
  |── e : C_execution_unit
```

If a single executor is defined in pss_top the hierarchy and exec order will be as follows:

```
pss_top
  |── c1 : my_comp
  |    |── action act
  |── c2 : my_comp
       |──action act
  |── e : C_execution_unit
  |── t : A_executor_c

1. run_start of pss_top.e
2. run_start of pss_top.t
```

Copyright © 2026 Accellera. All rights reserved.

```
3. run_start of pss_top.c1
4. run_start of pss_top.c2
5. run_start of action c1::act
6. run_start of action c2::act
```

20.2 Functions

Functions are a means to encapsulate behaviors used by **action**s and other entities to implement test scenarios. Functions are called in procedural description contexts, and are akin to procedures in conventional programming languages.

Functions can be declared in global, **package**, or **component** scopes. Functions can be *static* or *instance* (*non-static*) functions. A global or package function is always static. A component function can be explicitly declared as **static**. If a component function is non-static, each function call is associated with a specific instance of that **component** type.

A function may be defined in one of three ways:

- Using native PSS procedural statements, possibly calling other functions (see 20.3).
- As bound to a procedural interface in a foreign programming language, such as a function in C/C++, or a function/task in SystemVerilog (see 20.4). This only applies to static functions; an instance function cannot be bound.
- As a target code template block (see 20.6).

The definition of a function in one of these three ways may be coupled with the function's initial declaration. The definition may also be provided without a preceding declaration. In this case, the definition serves also as a declaration. On the other hand, the definition may be provided separately from the declaration, in a different lexical scope. The intent and semantics of a function are fixed by its declaration, but its implementation could vary between different environments and contexts. However, a given PSS model, along with all its source units, shall only contain one definition for any function (in case of an instance function, at most one definition per derived component type).

Functions may be called from procedural *exec* blocks, namely **exec init_down**, **init_up**, **pre_solve**, **post_solve**, **body**, **run_start**, and **run_end**. Functions called from **exec init_down**, **init_up**, **pre_solve**, and **post_solve** are evaluated on the *solve platform*, whereas functions called from **exec body**, **run_start** and **run_end** are evaluated on the *target platform*. If a function is explicitly declared as **solve** or **target**, its usage is restricted to the context of exec blocks of corresponding kinds.

A static function declared in a component scope may be shadowed by a function declaration with the same name in a derived component, which can be static or non-static. The function declaration in the derived component may have a different return type or arguments than in the base component.

An instance function declared in a component scope may be shadowed by an instance function declaration with the same name in a derived component. The function declaration in the derived component shall have the same return type and arguments as that in the base component. The function in the base type may be called from within the function in the derived type by calling "**super.**`<function name>(...)`".

However, an instance function cannot be shadowed by a static function.

When the shadowed element is an instance function, the function call is *polymorphic*, that is, the actual function called depends on its context component. See 17.1 for details. On the other hand, static functions calls are not polymorphic.

Copyright © 2026 Accellera. All rights reserved.

20.2.1 Function declarations

A function prototype is declared in a **package** or **component** scope within a PSS description. The function prototype specifies whether the function availability is restricted to a **solve** or **target** platform, whether it is static, the function name, return type, and function parameters. See Syntax 93. Note that the syntax shown here is for the declaration of a function prototype only, where the definition is provided separately. A function can also be declared and defined at once using a procedural statement block, a target code template, or an import function statement (see 20.3, 20.6, and 20.4, respectively). The same syntax is used for specifying the prototype in these cases also.

20.2.1.1 Syntax

```
function_decl ::=
    [ platform_qualifier ] [ pure ] [ static ] function function_prototype ;
platform_qualifier ::=
     target
   | solve
function_prototype ::=
   function_return_type function_identifier function_parameter_list_prototype
function_return_type ::=
    void
  | data_type
function_parameter_list_prototype ::=
    ( [ function_parameter { , function_parameter } ] )
  | ( { function_parameter , } varargs_parameter )
function_parameter ::=
    [ function_parameter_dir | const ] data_type identifier [ = constant_expression ]
  | [const] ( type | ref ref_type_category | plain_data_type ) identifier
function_parameter_dir ::=
    input
  | output
  | inout
varargs_parameter ::= ( data_type | type | ref ref_type_category | plain_type_category ) ... identifier
ref_type_category ::=
   action
 | monitor
 | component
 | object_kind
plain_type_category ::=
    struct
  | numeric
```

Syntax 93—Function declaration

The following also apply:

Copyright © 2026 Accellera. All rights reserved.

a) Functions declared in global or package scopes are considered static, regardless of whether the static qualifier is used.
b) The optional *platform_qualifier* (either **solve** or **target**) specifies function availability. An unqualified function is assumed to be available during all phases of test generation and execution.
c) Static functions (declared any scope) are called optionally using package or component type qualification with the scope operator (**::**).
d) Instance functions are called optionally using the dot operator (.) on a component instance expression.

20.2.1.2 Example

For an example of declaring a function, see 20.2.2, below.

20.2.1.3 Specifying function availability

In some environments, test generation and execution are separate activities. In those environments, some functions may only be available during test generation, on the *solve platform*, while others are only available during test execution, on the *target platform*. For example, reference model functions may only be available during test generation while the utility functions that program hardware devices may only be available during test execution.

An unqualified function is assumed to be available during all phases of test generation and execution. Qualifiers are specified to restrict a function's availability. Functions restricted to the *solve platform* shall not be called directly or indirectly from *target exec*s, namely **body**, **run_start**, and **run_end**. Similarly, functions restricted to the *target platform* shall not be called from *solve exec*s, namely **init_down**, **init_up**, **pre_solve**, **post_solve**, and **pre_body**.

Example 288 specifies function availability. Two functions are declared in the `external_functions_pkg` package. The `alloc_addr` function allocates a block of memory, while the `transfer_mem` function causes data to be transferred. Both of these functions may be present in all phases of test execution in a system where solving is done on-the-fly as the test executes; therefore, no platform qualifier is used.

In a system where a pre-generated test is to be compiled and run on an embedded processor, memory allocation may be pre-computed. Data transfer shall be performed when the test executes. The `pregen_tests_pkg` package specifies these restrictions: `alloc_addr` is only available during the solving phase of stimulus generation, while `transfer_mem` is only available during the execution phase of stimulus generation. PSS processing uses this specification to ensure that the way imported functions are used aligns with the restrictions of the target environment.

Copyright © 2026 Accellera. All rights reserved.

```
package external_functions_pkg {

  function bit[32] alloc_addr(bit[32] size);

  function void transfer_mem(
    bit[32] src, bit[32] dst, bit[32] size
  );
}

package pregen_tests_pkg {

  import solve function external_functions_pkg::alloc_addr;

  import target function external_functions_pkg::transfer_mem;

}
```

Example 288—Function availability

Example 289 demonstrates an **activity** with reactive control flow based on values returned from a target function called in an **exec body** block.

```
component my_ip_c {
  import target C static function int sample_DUT_state();
      // specify mapping to target C function by that same name
  action check_state {
    int curr_val;
    exec body {
      curr_val = sample_DUT_state();
          // value only known during execution on target platform
    }
  };
  action A { };
  action B { };
  action my_test {
    check_state cs;
    activity {
      repeat {
        cs;
        if (cs.curr_val % 2 == 0) {
          do A;
        } else {
          do B;
        }
      } while (cs.curr_val < 10);
    }
  };
};
```

Example 289—Reactive control flow

20.2.2 Parameters and return types

A function shall explicitly specify a data type as its return type or use the keyword **void** to indicate that the function does not return a value. Function return values shall be either plain-data types (scalars and

Copyright © 2026 Accellera. All rights reserved.

aggregates thereof) or reference types. Functions shall not return **action** types, **component** types, or flow/resource object types without the **ref** modifier.

A function may specify any number of formal parameters, stating their types and names. Function parameters shall be either plain-data types or reference types. Functions shall not have parameters of **action** types, **component** types, or flow/resource object types without the **ref** modifier. Functions may also declare generic parameters without stating their specific type, and may declare a variable number of parameters—see 20.2.5. Note that the set of types allowed for imported foreign functions is restricted (see 20.4).

Parameter direction modifiers (**input**, **output**, or **inout**) are optional in the function declaration. However, if they are specified in the function declaration, such a function may only be imported (see 20.4). Functions whose definition is built into implementations, such as functions included in the PSS core library (see Clause 21), may also have **output** or **inout** parameters. In the declaration of native functions and target-template functions, direction modifiers shall not be used.

Example 290 declares a function in a **package** scope. In this case, the function `compute_value` returns an **int**, accepts an input value (`val`), and returns an output value via the `out_val` parameter.

```
package generic_functions {
  function int compute_value(
    int        val,
    output int out_val);
}
```

Example 290—Function declaration

20.2.3 Const parameters

const is an optional qualifier that may be associated with one of the function parameters.

A function parameter declared with the *const* qualifier is considered constant in the scope of the function, i.e., this parameter value cannot be changed in the function body. A constant of an aggregate data type can be passed as an argument only to functions in which the corresponding parameter has a *const* qualifier.

An aggregate literal passed as an argument is a constant in the context of the function call, regardless of whether the fields of the aggregate literal are themselves constant expressions. Therefore, the argument in the function prototype requires the *const* qualifier.

The following also apply:

a) It shall be illegal to provide both *const* and *function_parameter_dir* qualifiers.
b) It shall not be allowed to add the *const* qualifier on reference types or collections thereof.
c) The *const* qualifier is an essential part of the function signature and shall appear in redeclarations (and overrides) of a function.

Example 291 demonstrates declarations and usage of const parameters.

Copyright © 2026 Accellera. All rights reserved.

```
struct s {
  int a, b;
}
static const array<int, 3> MY_ARR = {1,2,3};
static const s MY_STRUCT = {.a = 1, .b =2};
function void copy_array(const array<int, 3> src, array<int, 3> dst) {
  foreach (dst[i]) {
    dst[i] = src[i];
    // src[i] = dst[i];    // ERROR - src is const and its value cannot
                           // be changed
  }
}
function void swap_structs(s s1, s s2) {
    s tmp = s1;
    s1 = s2;
    s2 = tmp;
}
component pss_top {
  array<int, 3> my_arr;
  s s1, s2;
  exec init_up {
    copy_array(my_arr, MY_ARR);   // ERROR: dst parameter is not const
    copy_array(my_arr, {1,2,3});  // ERROR: dst parameter is not const
    copy_array(MY_ARR, my_arr);   // OK
    swap_structs(s1, MY_STRUCT);  // ERROR: s2 parameter is not const
    swap_structs(s1, s2);         // OK
  }
}
```

Example 291—Const parameter declaration

In Example 291 above, the function `copy_array` accepts two array parameters `src` and `dst`. `src` is declared with the const qualifier; therefore, inside the function body, its elements cannot be modified. `dst` is not declared as const. The first call of `copy_array` in the exec `init_up` is illegal because `MY_ARR` is a static constant but passed as the second argument, the non-const `dst` parameter. Similarly, the second call of `copy_array` in the exec `init_up` is illegal because the literal expression `{1,2,3}` is treated as a constant in the context of the function call.

The first call to the `swap_structs` function is illegal because the static constant struct `MY_STRUCT` is passed as an argument to a function whose parameters are declared non-const.

20.2.4 Default parameter values

Default parameter values serve as the actual values for the respective parameters if explicit actual parameters are missing in the function call.

The following also apply:

a) A default parameter value shall be specified as a constant expression, and therefore can only be specified for a parameter of a plain-data type.

b) In a function declaration, following a parameter with a specified default value, all subsequent parameters shall also have default values specified.

c) A default parameter value is in effect for redeclarations (and overrides) of a function. The default parameter value may be specified in the redeclaration of a function if already declared for the same parameter in a previous declaration, but this value shall be equal to the original one.

Copyright © 2026 Accellera. All rights reserved.

d) In an import function declaration, default parameters are not allowed on **output** or **inout** arguments.

Example 292 demonstrates the declaration and use of a default parameter value.

```
function void foo(int x, int y = 100);
function void bar() {
   foo(3,200); // the value 200 is used for parameter y
   foo(3);     // the value 100 is used for parameter y
}
```

Example 292—Default parameter value

20.2.5 Generic and varargs parameters

Generic parameters and *varargs parameters* are means to declare functions that are generic or variadic with respect to their parameters. Examples are functions that apply to all actions or objects as such, and functions that involve string formatting.

Generic and varargs parameters are used for the declaration of functions whose definition is built into implementations. In particular, they are used to declare functions included in the PSS core library (see Clause 21). PSS does not provide a native mechanism to operate on an unspecified number of parameters or on parameters with no declared type, nor does PSS define mapping of functions with generic/varargs parameters to foreign languages.

The following also apply:

a) A generic parameter is declared either with the keyword **type** or with a *type category*, rather than with a specific type. A value of any type (if **type** was specified), or any type that belongs to the specified category (if a type category was specified), is accepted in the function call. In the case of the **struct** or **numeric** category, the **ref** modifier shall not be used, but for the other categories (**component**, **action**, one of the object kinds), the **ref** modifier shall be used.

b) Default values may not be specified for generic parameters.

c) The varargs parameter (ellipsis notation – "`...`") signifies that zero or more trailing values may be passed as actual parameters in the function call. Note that a varargs parameter may only occur as the last parameter in the parameter list.

d) In a function call, the expressions corresponding to a varargs parameter shall all be of the declared type if a type is specified, or belong to the same type category if one is specified. Note that in the case of a type category, the types of the actual parameter expressions may vary, so long as they all belong to the specified category. When a varags parameter is declared with the keyword **type**, actual parameters types may vary with no restriction.

Example 293 demonstrates the declaration and use of a generic parameter.

```
function void foo(struct x);
struct my_struct {};
struct your_struct {};
function void bar() {
   my_struct s1;
   your_struct s2;
   foo(s1);
   foo(s2);
}
```

Example 293—Generic parameter

Copyright © 2026 Accellera. All rights reserved.

Example 294 demonstrates the declaration and use of a varargs parameter.

```
function string format_string(string format, type ... args);
function void bar() {
   string name = "John";
   int age = 55;
   string result;
   result = format_string("name %s: age %d", name, age);
}
```

Example 294—Varargs parameter

20.2.6 Pure functions

Pure functions are functions for which the return value depends only on the values of their parameters, and their evaluation has no side-effects. Declaring a function as **pure** may provide the PSS implementation with opportunities for optimization. Note that a function declared as **pure** may lead to unexpected behavior if it fails to obey these rules.

The following rules apply to **pure** functions, that is, functions declared with the **pure** modifier:

a) Only non-void functions with no **output** or **inout** parameters may be declared **pure**.

b) The **pure** keyword may be omitted in a function definition if its original declaration contains the **pure** keyword; it is still considered pure.

A non-**pure** function shall not be declared as **pure** in derived types.

20.2.6.1 Example

Example 295 demonstrates declaration and use of **pure** functions.

```
pure function int factorial(int n);
action A {
   rand int vals[10];
   int factorial_vals[10];

   exec post_solve {
      foreach (vals[i]) {
         factorial_vals[i] = factorial(vals[i]);
      }
   }
}
```

Example 295—Pure function

In the example above, the function `factorial()` is pure and therefore will not necessarily be re-evaluated for each element in the array. If some elements in the array are equal, the PSS implementation may choose to use the result of a previous evaluation, and not evaluate the function again.

20.2.7 Calling functions

Functions may be called directly from *exec block*s or from other functions using *procedural constructs* (see 20.7). Recursive function calls are allowed.

Copyright © 2026 Accellera. All rights reserved.

Functions not returning a value (declared with **void** return type) may only be called as standalone procedural statements. Functions returning a value may be used as operands in expressions; the value of that operand is the value returned by the function. The function can be used as a standalone statement and the return value discarded by casting the function call to **void**:

```
(void)function_call();
```

Calling a nonvoid function as if has no return value shall be legal, but it is recommended to explicitly discard the return value by casting the function call to **void**, as shown above.

Example 296 demonstrates calling various functions. In this example, the `mem_segment_s` **buffer** object captures information about a memory buffer with a random size. The specific address in an instance of the `mem_segment_s` object is computed using the `alloc_addr` function. `alloc_addr` is called after the solver has selected random values for the **rand** fields (specifically, `size` in this case) to select a specific address for the `addr` field.

```
package external_functions_pkg {
  function bit[32] alloc_addr(bit[32] size);

  function void transfer_mem(
    bit[32] src, bit[32] dst, bit[32] size
  );

  buffer mem_segment_s {
    rand bit[32]        size;
    bit[32]             addr;

    constraint size in [8..4096];

    exec post_solve {
      addr = alloc_addr(size);
    }
  }
}

component mem_xfer {
  import external_functions_pkg::*;

  action xfer_a {
    input  mem_segment_s     in_buff;
    output mem_segment_s    out_buff;

    constraint in_buff.size == out_buff.size;

    exec body {
      transfer_mem(in_buff.addr, out_buff.addr, in_buff.size);
    }
  }
}
```

Example 296—Calling functions

A function call shall only be valid if the function has an existing definition available on the relevant platform. In case of an instance function, a function call shall be valid if the function has an existing definition available on the relevant platform, defined in context of the relevant component instance type. It is

Copyright © 2026 Accellera. All rights reserved.

possible to declare an instance function in context of a component base type but provide a definition only in context of its subtypes, as long as no function call is done in context of a base type instance.

Example 297 demonstrates the function call restrictions when a function is declared but not defined. It declares various functions in the global package or in a component, some of them restricted to a **solve** or **target** platform. Then various definitions are provided to these functions, some add a platform restriction not specified in the original declaration. Some of the function calls within the **post_solve** and **body** exec blocks are valid, some are invalid, according to these restrictions.

Copyright © 2026 Accellera. All rights reserved.

```
import std_pkg::*;

function int func1(int a, int b);
function int func2(int a, int b);
solve function int func3(int a);
target function void func4(int a);

function int func1(int a, int b) {  // definition for any platform
    return a+b;
}
solve function int func2(int a, int b) {  // solve-only definition
    return a+b;
}
function int func3(int a) {  // definition for the solve function
    int res;
    randomize res with {
        res > a;
    };
    return res;
}
function void func4(int a) {  // definition for the target function
    message(LOW, "The value is %d", a);
}

component base_comp {
    function void func5();

    action do_it {
        int a, b, c, x, y, z;
        exec post_solve {
            x = func1(1,2);
            y = func2(3,4);
            z = func3(5);
            func4(6);     // ERROR: calling a target function
            comp.func5(); // OK under comp1, ERROR under comp2
        }
        exec body {
            a = func1(1,2);
            b = func2(3,4); // ERROR: calling a solve function
            c = func3(5);   // ERROR: calling a function
                            // with a solve-only definition
            func4(6);
            comp.func5();   // OK under comp2, ERROR under comp1
        }
    }
}
component comp1: base_comp {
    solve function void func5() {  // solve-only definition
        print("I am comp1::func5");
    }
}
component comp2: base_comp {
    target function void func5() {  // target-only definition
        message(LOW, "I am comp2::func5");
    }
}
```

Example 297—Function calls restrictions

Copyright © 2026 Accellera. All rights reserved.

20.3 Native PSS functions

It is possible to specify the definition for native PSS functions using the procedural constructs described in 20.7.

For an instance function, the definition (if provided) shall be in the same **component** type as the original declaration (either in its initial definition or in an extension) or in a derived **component**. For a static function, the definition shall be in the same **package** or **component** as the original declaration (in case of a component, either in its initial definition or in an extension).

20.3.1 Syntax

```
procedural_function ::= [ platform_qualifier ] [ pure ] [ static ] function
    function_prototype { { procedural_stmt } }
platform_qualifier ::=
    target
  | solve
function_prototype ::= function_return_type function_identifier function_parameter_list_prototype
function_return_type ::=
    void
  | data_type
function_parameter_list_prototype ::=
    ( [ function_parameter { , function_parameter } ] )
  | ( { function_parameter , } varargs_parameter )
function_parameter ::=
    [ function_parameter_dir | const] data_type identifier [ = constant_expression ]
  | [const] ( type | ref ref_type_category | plain_type_category ) identifier
function_parameter_dir ::=
    input
  | output
  | inout
varargs_parameter ::= ( data_type | type | ref ref_type_category | plain_type_category ) ...
    identifier
ref_type_category ::=
    action
  | monitor
  | component
  | object_kind
plain_type_category ::=
    struct
  | numeric
```

Syntax 94—Function definition

The optional *platform_qualifier* (either **solve** or **target**) specifies function availability. If the function declaration is provided separately and is qualified, *platform_qualifier* shall be the same, or it may be

Copyright © 2026 Accellera. All rights reserved.

omitted. If the function declaration is unqualified, *platform_qualifier* restricts the function availability for the PSS model. If no separate function declaration exists, this definition also serves as the declaration, and *platform_qualifier* or its absence is treated accordingly (see 20.2.1 and Example 297).

For native PSS functions, *function_parameter_dir* shall be left unspecified for all parameters of the function, both in the original function declaration (if provided) and in the native PSS function definition.

20.3.2 Parameter passing semantics

Parameter direction shall be unspecified in the function prototype for native PSS functions. This implies that the parameter direction (**input**, **output**, or **inout**) shall not be used. If the function declaration contains directions for parameters, this function shall not have a native implementation.

In the implementation of these functions, the following apply:

- Parameters of scalar data types are passed by value. Any changes to these parameters in the callee do not update the values in the caller.
- Parameters of aggregate data types are passed as a handle to the instance in the caller. Updates to these parameters in the callee will modify the instances in the caller. When a variable of inherited type is passed as a parameter of base type, only the fields present in the base type are visible within the function. Note that as variables, parameters of aggregate data types have value semantics in assignment and equality expressions (see 8.3 and 8.5.3).
- Parameters of reference data types are passed as reference assignments. The parameter points to (is an alias to) the entity referred to in the actual parameter expression. Note that as variables, parameters of reference types have reference semantics in assignment and equality expressions (see 8.3 and 8.5.3), and may evaluate to **null**.

Example 298 shows the parameter passing semantics.

Copyright © 2026 Accellera. All rights reserved.

```
package generic_functions {
   struct params_s {
      int x;
   };

   struct params_inh_s : params_s {
      int y;
   }

   // Prototypes
   function void set_val0(params_s p, int a);
   function void set_val1(params_s p_dst, params_s p_src);
   function params_s zero_attributes();

   // Definitions
   function void set_val0(params_s p, int a)
   {
      p.x = a;
      a = 0;
   }
   function void set_val1(params_s p_dst, params_s p_src)
   {
      p_dst.x = p_src.x;
   }
   function params_s zero_attributes()
   {
      params_s s;
      s.x = 0;
      return s;
   }

   component A {
      params_s p;
      params_inh_s p_inh1, p_inh2;
      int a;

      exec init_up {
         a = 10;
         p.x = 20;
         set_val0(p, a);
         // p.x is set to 10 at this point and a is unchanged

         set_val1(p, zero_attributes());
         // p.x is set to 0 at this point

         // Variables of inherited type may be passed as
         // function parameters of base type
         p_inh1.x = 5;
         p_inh1.y = 15;
         p_inh2.x = 10;
         p_inh2.y = 20;
         set_val1(p_inh1, p_inh2);
         // The value of p_inh1.y can never be changed by set_val1 because
         // set_val1 can only access fields of params_s (i.e., x)
      }
   };
}
```

Example 298—Parameter passing semantics

Copyright © 2026 Accellera. All rights reserved.

20.4 Foreign procedural interface

PSS defines a foreign procedural interface to enable calling imported external functions and interface-class members from PSS and enable calling exported PSS functions from the environment.

20.4.1 Definition using imported functions

Static function declarations in PSS may expose, and ultimately be bound to, foreign language APIs (functions, tasks, procedures, etc.) available on the target platform and/or on the solve platform. A function that was previously declared in the PSS description can be designated as *imported*. Calling an imported function from a PSS procedural context invokes the respective API in the foreign language. Parameters and result passing are subject to the type mapping defined for that language.

Instance functions cannot be imported.

Additional language qualifiers are included with imported functions to provide more information to the tool about the way the function is implemented. In typical use, such qualifiers are specified in an environment-specific package (e.g., a UVM environment-specific package or C-test-specific package).

A **static** function declared in a **component** can only be imported in the scope of the same component type. It shall be illegal to import a function declared in a base component type within a derived or unrelated component type.

It shall be illegal to import a function declared in a *template* **component** type.

20.4.1.1 Syntax

```
import_function ::=
    import [ platform_qualifier ] [ language_identifier ] function type_identifier ;
  | import [ platform_qualifier ] [ language_identifier ] [ static ] function function_prototype ;
platform_qualifier ::=
    target
  | solve
```

Syntax 95—Imported function qualifiers

The following also apply:

a) The first form of *import_function* can only be used when a separate function declaration is provided.

b) The optional *platform_qualifier* (either **solve** or **target**) specifies function availability. If the function declaration is provided separately and is qualified, *platform_qualifier* shall be the same, or it may be omitted. If the function declaration is unqualified, *platform_qualifier* restricts the function availability for the current PSS model. If no separate function declaration exists, this definition also serves as the declaration, and *platform_qualifier* or its absence is treated accordingly (see 20.2.1 and Example 297).

c) Return values and parameter values of imported functions are restricted to the following types:

   1) **bit** or **int**, provided width is no more than 64 bits
   2) **bool**
   3) **enum**
   4) **string**

Copyright © 2026 Accellera. All rights reserved.

5) **chandle**
6) **struct**
7) **array** whose element type is one of those listed in 1-6 above, including a sub-array
8) **list** whose element type is one of those listed in 1-6 above, except **string** or **chandle**

See Annex D for type-mapping rules to C, C++, and SystemVerilog.

d) Parameter direction modifiers may be used in the **function** declaration or in the **import** declaration to specify the passing semantics between PSS and the foreign language:
   1) If the value of an **input** parameter is modified by the foreign language implementation, the updated value is not reflected back to the PSS model.
   2) An **output** parameter sets the value of a PSS model variable. The foreign language implementation shall consider the value of an **output** parameter to be unknown on entry; it shall specify a value for an **output** parameter.
   3) An **inout** parameter takes an initial value from a variable in the PSS model and reflects the value specified by the foreign language implementation back to the PSS model.
e) In the absence of an explicit direction modifier, parameters default to **input**.

In addition, the following apply when the second form of *import_function* is used (with the function prototype specified):

a) If the direction for a parameter is left unspecified in the **import** declaration, it defaults to **input**.
b) The prototype specified in the **import** declaration shall match the prototype specified in the **function** declaration in the following ways:
   1) For a **static** function declared in a **component**, the **static** qualifier shall be used.
   2) The number of parameters shall be identical.
   3) The parameter names, types, and directions shall be identical.
   4) The return types shall be identical.

20.4.1.2 Specifying an implementation language

The implementation language for an imported function can be specified implicitly or explicitly. In many cases, the implementation language need not be explicitly specified because the PSS processing tool can use sensible defaults (e.g., all imported functions are implemented in C++). Explicitly specifying the implementation language using a separate statement allows different imported functions to be implemented in different languages, however (e.g., reference model functions are implemented in C++, while functions to drive stimulus are implemented in SystemVerilog).

Example 299 shows explicit specification of the foreign language in which the imported function is implemented. In this case, the function is implemented in C. Notice that only the name of the imported function is specified and not the full function prototype.

```
package known_c_functions {
    import C function generic_functions::compute_expected_value;
}
```

Example 299—Explicit specification of the implementation language

20.4.2 Definition using exported functions

Exported functions allow the environment to call back into the implementation of a PSS model. The export qualifier may be associated with a static PSS function having a native implementation (i.e., not an import or

Copyright © 2026 Accellera. All rights reserved.

target-template implementation). Only static functions defined in a component or in a package may be exported.

20.4.2.1 Syntax

```
export_function ::= export target function type_identifier ;
```

Syntax 96—Exported function

Return values and parameter values of exported functions are restricted to the following types:

a) **bit** or **int**, provided width is no more than 64 bits
b) **bool**
c) **enum**, provided width of the base type is no more than 64 bits
d) **float32** or **float64**
e) string
f) **chandle**
g) **struct**, provided all fields are of a type supported for export functions and provided that the function declares the parameter as “const”

See Annex D for type-mapping rules to C, C++, and SystemVerilog.

Copyright © 2026 Accellera. All rights reserved.

20.4.2.2 Examples

```
import sync_pkg::*;
import executor_pkg::*;

// Hands the environment a handle naming the executor PSS is running on
import target function void irq_register(chandle ctxt);

// Environment APIs on the target platform
import target function void    dma_program(bit[8] chan, bit[64] src,
                                           bit[64] dst, bit[32] len);
import target function void    dma_start(bit[8] chan);
import target function bit[32] dma_ack_irq(bit[8] chan);
// read and clear

component dma_c {
  static const int NCHAN = 4;

  // One rendezvous per hardware channel, so concurrent transfers
  // don't race for each other's completions
  channel_c<bit[32]> completion[NCHAN];

  // What the C handler calls back into. Runs in interrupt context,
  // so it uses try_put: a full channel must not become a
  // hung handler.

  action xfer_a {
    rand bit[8]  chan;
    rand bit[64] src;
    rand bit[64] dst;
    rand bit[32] len;

    constraint chan < NCHAN;

    exec body {
      dma_program(chan, src, dst, len);
      dma_start(chan);

      // Blocks here. No polling loop, no yield, no
      // status register read -- the handler will deliver the
      // status when the DMA finishes.
      bit[32] status = comp.completion[chan].get();

      if (status != 0) {
        error("DMA channel %d completed with status %d", chan, status);
      }
    }
  }
}
```

Example 300—Interrupt-driven PSS DMA model

Copyright © 2026 Accellera. All rights reserved.

```
component pss_top {
  dma_c dma;

  // Give the environment the context handle before any test code runs
  exec run_start {
    irq_register(executor().get_context());
  }
}

function void dma_isr(bit[8] c) {
  pss_top.dma.completion[c].try_put(dma_ack_irq(c));
}
export target function dma_isr;
```

Example 300—Interrupt-driven PSS DMA model (cont.)

Example 300 shows the PSS portion of an interrupt-driven model. The DMA component uses a per-DMA-channel `channel_c` instance to enable waiting for transfer completion. The `dma_isr` function is exported by the PSS model, allowing the environment to notify the model when an interrupt occurs.

```
/* Generated by the PSS tool for 'export target function dma_isr'.
 The leading context parameter is what tells the runtime which
 executor is calling. */
extern void dma_isr(void *ctxt, unsigned char chan);

static void *irq_ctxt;

/* Called from the PSS model's run_start exec */
void irq_register(void *ctxt)
{
  irq_ctxt = ctxt;
  plic_enable(IRQ_DMA);
}

/* Ordinary platform interrupt vector */
void sys_isr(int code)
{
  switch (code) {
    case DMA: dma_isr(irq_ctxt, dma_pending_chan()); break;
    // ...
  }
}
```

Example 301—Environment interrupt-handling functions

Example 301 shows the environment portion of interrupt handling. The `irq_register` function is called when the PSS test scenario starts to register a context handle that is passed back to the PSS test scenario when an interrupt occurs. The `sys_isr` function is registered with the system as an interrupt handler, and it is invoked when a system interrupt occurs. When the interrupt is due to the DMA engine, the handler calls the exported PSS function to handle the interrupt.

Copyright © 2026 Accellera. All rights reserved.

20.4.3 Imported classes

In addition to interfacing with external foreign language functions, the PSS description can interface with foreign language classes. See also Syntax 97.

20.4.3.1 Syntax

```
import_class_decl ::= import class import_class_identifier [ import_class_extends ]
  { { import_class_function_decl } }
import_class_extends ::= : type_identifier { , type_identifier }
import_class_function_decl ::= function_prototype ;
```

Syntax 97—Import class declaration

The following also apply:

a) Imported class functions support the same return and parameter types as imported functions. **import class** declarations also support capturing the class hierarchy of the foreign language classes.

b) Fields of **import class** type can be instantiated in **package** and **component** scopes. An **import class** field in a **package** scope is a global instance. A unique instance of an **import class** field in a **component** exists for each component instance.

c) Imported class functions are called from an *exec block* just as imported functions are.

20.4.3.2 Example

Example 302 declares two imported classes. **import class** `base` declares a function `base_function`, while **import class** `ext` extends from **import class** `base` and adds a function named `ext_function`.

```
import class base {
    void base_function();
}

import class ext : base {
    void ext_function();
}
```

Example 302—Import class

20.5 Target-template implementation of exec blocks

Implementation of **exec**s may be specified using a *target template*—a string literal containing code in a specific foreign language, optionally embedding references to fields in the PSS description. Target-template implementation is restricted to *target exec* kinds (**body**, **run_start**, **run_end**, **header**, and **declaration)**. In addition, target templates can be used to generate other text files using **exec file**. Target-template implementations may not be used for *solve exec*s **(init_down**, **init_up**, **pre_solve**, **post_solve**, and **pre_body**).

Target-template **exec**s are inserted by the PSS tool verbatim into the generated test code, except when a triple-quoted string with special elements is used (see 4.7.1). In the latter case, mustache expressions are expanded, control-flow directives are handled accordingly to produce parts of the generated code, and embedded comments are discarded. Multiple target-template *exec block*s of the same kind are allowed for a

Copyright © 2026 Accellera. All rights reserved.

given action, flow/resource object, or **struct**. They are (logically) concatenated in the target file, as if they were all concatenated in the PSS source.

20.5.1 Target language

A *language_identifier* serves to specify the intended target programming language of the code block. Clearly, a tool supporting PSS shall be aware of the target language to implement the runtime semantics. PSS does not enforce any specific target language support, but recommends implementations reserve the identifiers **`C`**, **`CPP`**, and **`SV`** to denote the languages C, C++, and SystemVerilog respectively. Other target languages may be supported by tools, given that the abstract runtime semantics are kept. PSS does not define any specific behavior if an unrecognized *language_identifier* is encountered.

Each target-template **exec** block is restricted to one target language in the context of a specific generated test. However, the same **action** may have target-template **exec** blocks in different languages under different **package**s, given that these **package**s are not used for the same test.

20.5.2 Exec file

Not all the artifacts needed for the implementation of tests are coded in a programming language that tools are expected to support as such. Tests may require scripts, command files, make files, data files, and files in other formats. The **exec file** construct (see 20.1) specifies text to be generated out to a given file. **exec file** constructs of different actions/objects with the same target are concatenated in the target file in their respective scenario flow order. The exec file ordering follows the ordering for the *run_start* exec blocks (See 20.1.5).

Note that a triple-quoted string including special elements such as mustache expressions (see 4.7.1) can be used to specify the filename to which exec file output is directed. This allows different traversals of the same action type to output content to different files.

```
action A {
    rand bit[4]        fileid;
    rand bit[4]        data;

    exec file """datafile_{{fileid}}""" = """data: {{data}}""";
}
```

Example 303—Mustache expression used to specify a filename

In Example 303, the output filename is specified in terms of a random action field. Data will be output to files named `datafile_0 .. datafile_15`, depending on the randomized value of `fileid`.

20.5.3 Controlling the output of target-template exec blocks

Implementing test intent requires using data from the PSS model in the code created from target-template exec blocks. This can be achieved by using a triple-quoted string with special elements, such as mustache expressions or control-flow directives (see 4.7.1). Within these special elements, expressions can be used involving variables declared in the scope in which the exec block is declared.

Example 304 shows referencing PSS variables inside a target-template exec block using mustache notation.

Copyright © 2026 Accellera. All rights reserved.

```
component top {
    struct S {
        rand int b;
    }
    action A {
        rand int a;
        rand S   s1;
        exec body C = """
          printf("a={{a}} s1.b={{s1.b}} a+b={{a+s1.b}}\n");
        """;
    }
}
```

Example 304—Referencing PSS variables using mustache notation

A variable reference can be used in any position in the generated code. Example 305 shows a variable reference used to select the function being called.

```
component top {
    action A {
        rand bit[2] func_id;
        rand bit[4] a;
        exec body C = """
          func_{{func_id}}({{a}});
        """;
    }
}
```

Example 305—Variable reference used to select the function

One implication of this is that a mustache reference cannot be used to assign a value to a PSS variable.

Example 305 also declares a random `func_id` variable that identifies a C function to call. When a PSS tool processes this description, the following output shall result, assuming `func_id==1` and `a==4`:

```
func_1(4);
```

Example 306 shows how a procedural **pre_solve** exec block is used along with a target-template declaration exec block to allow programmatic declaration of target variable declarations. Besides mustache expression, the target-template block in this example uses a control-flow directive **foreach**.

Copyright © 2026 Accellera. All rights reserved.

```
enum obj_type_e {my_int8,my_int16,my_int32,my_int64};
import solve function string get_unique_obj_name();

buffer mem_buff_s {
  rand array<obj_type_e,3> obj_types;
  array<string,3> obj_names;

  exec post_solve {
    repeat (i: 3)
      obj_name[i] = get_unique_obj_name();
  }

  // declare an object in global space
  exec declaration C = """
    {% foreach (obj_name: obj_names[i]) %}
      static {{obj_types[i]}} {{obj_name}};
    {%%}
  """;
}
```

Example 306—Allowing programmatic declaration of target variable declarations

Assume that the solver selects `my_int16`, `my_int64`, and again `my_int16` as the values of the `obj_types` array and that the `get_unique_obj_name()` function returns `field__0`, `field__1`, and `field__2` in this order. In this case, the PSS processing tool shall generate the following content in the declaration section:

```
static my_int16 field__0;
static my_int64 field__1;
static my_int16 field__2;
```

20.5.4 Exec block tag

target_code_exec_block ::= **exec** exec_kind *language*_identifier **=**
[exec_block_tag **:**] string_literal **;**
target_file_exec_block ::= **exec file** filename_string **=** [exec_block_tag **:**] string_literal **;**
exec_block_tag ::= type_identifier [struct_literal]
filename_string ::= string_literal

Syntax 98—Exec block tag

A target template exec block may be associated with a tag value used to determine equivalence between exec blocks during generation. The tag is used solely for determining whether multiple exec blocks match for the purposes of code generation deduplication. The exec tag does not affect action traversal, solve behavior, or runtime execution.

The exec tag may be specified for target template exec blocks of kinds *header*, *declaration*, *run_start*, *run_end*, and exec *file* blocks. The tag is not permitted on native exec blocks, on *body* (which is invoked per traversal), or on solve execs (*init_down*, *init_up*, *pre_solve*, *post_solve*, *pre_body*).

The exec tag shall be specified as a struct type, optionally followed by a struct literal. The exec tag shall be evaluated using a temporary instance of the specified struct type, constructed at the time the target template

Copyright © 2026 Accellera. All rights reserved.

exec block is evaluated (after execution of exec pre_body). The instance shall be initialized using the specified struct literal, if present; otherwise, by using the field default values. The instance shall not participate in randomization, and no exec blocks declared within the struct type shall be evaluated as part of tag evaluation.

During code generation, two exec blocks of the same kind whose tags are equal are coalesced except if this is the same instance.

20.5.4.1 Matching rules

Two exec blocks are considered to have the same tag value if their tags are of the same struct type and their evaluated values are equal according to the equality rules defined in 8.5.3. Tag values of different struct types shall not match. Exec blocks that do not specify an exec tag shall not match any other exec block, including another untagged exec block.

Two exec blocks match if:

a) they have the same *exec_kind*
b) they have the same *language_identifier*
c) their executors are associated with the same target execution unit instance
d) their tag expressions have the same struct type
e) their evaluated tag values are equal

Two exec *file* blocks match if:

a) their filename matches
b) their tag expressions have the same struct type
c) their evaluated tag values are equal

Matching is independent of declaration scope. Exec blocks declared in different types, including different action types or components, may match.

Within the same instantiated object, multiple matching exec blocks shall be concatenated in declaration order (see Example 310).

Exec blocks are evaluated as specified in the exec ordering rules (see 20.1.5). When a tagged exec block is generated and a previously generated matching exec block originating from a different instantiated object exists, the generated text shall be compared with the previously generated text for that match. If the generated text is identical on a character-by-character basis, the current exec block shall be suppressed. If the generated text is not identical on a character-by-character basis, an error shall be reported.

20.5.4.2 Exec file tag

Untagged exec file blocks do not match and their contents are concatenated. Tagged exec file blocks participate in tag based coalescing. When multiple tagged exec file blocks match, only one contributes content.

Tagged and untagged contributions to the same file coexist, and the final file content is the concatenation of all contributions in source order.

Copyright © 2026 Accellera. All rights reserved.

20.5.4.3 Examples

```
struct shared_helpers_tag {}
struct shared_helpers_s {
    exec declaration C = shared_helpers_tag : """
        static int helper_count = 0;
        static void shared_helper(int x) { helper_count += x; }
    """;
};
action a1 { shared_helpers_s h; };
action a2 { shared_helpers_s h; };
```

Example 307—Constant tag—generated once across instances

In Example 307, a scenario traversing `a1` and `a2` any number of times generates the declaration block exactly once.

```
enum cache_op_e {CLEAN, INVALIDATE, CLEAN_INVALIDATE};
struct cache_op_tag {
    cache_op_e kind;
}
action cache_op {
    rand cache_op_e kind;
    exec header C = cache_op_tag{.kind = kind} : """
       static inline void cache_op_{{kind}}(uint64_t addr) { /* ... */ }
    """;

    exec body C = "cache_op_{{kind}}(addr);";
}
```

Example 308—Evaluated tag—generated per attribute value

In Example 308, a scenario with three traversals resolving to kind `{CLEAN, CLEAN, INVALIDATE}` generates the header exactly twice—once for `CLEAN` and once for `INVALIDATE`. The untagged body is generated at every call and invokes the appropriate helper function introduced by the deduplicated header.

```
component cpu_executor_c : executor_c<> {
    struct boot_tag {
        string target_id;
    }
    string target_id;

    exec file """boot_{{target_id}}.S""" = boot_tag{.target_id =
                        target_id} : """
        .global _start_{{target_id}}
        _start_{{target_id}}: /* ... */
    """;
}
```

Example 309—Tagged exec file in an executor component

Copyright © 2026 Accellera. All rights reserved.

In Example 309, when two or more instances of `cpu_executor_c` resolve to the same `target_id` (e.g., several actions targeting the same CPU), the boot stub is written into `boot_<target_id>.S` exactly once. Distinct `target_id` values produce distinct files.

```
component my_comp {
  struct tag {
    string name;
  }
  action A {
    exec declaration C = tag{.name="A"} : """
    // Action A
    """;
  }
  action B : A {
    exec declaration C = tag{.name="B"} : """
    // Action B
    """;
  }
  action C : A {
  }
  action D : A {
    exec declaration C = tag{.name="A"} : """
    // Action D
    """;
  }
  extend action A {
    exec declaration C = tag{.name="A"} : """
    // Action A - extension
    """;
  }
  action test {
    activity {
      do A;
      do B;
      do C;
      do D;
    }
  }
}
```

Example 310—Exec tag with extension and inheritance

Example 310 highlights the interaction between inheritance, extension, shadowing, and tag-based coalescing: extension contributes additional exec blocks within the same type, inheritance brings them into derived actions, shadowing removes base declarations when a derived action defines its own, and tag matching then determines whether blocks are emitted, suppressed, or cause errors across different action instances.

For action `A`, the execs in the original definition and the extension share the same tag, so they are concatenated and both generated (no suppression within the same instance):

```
// Action A

// Action A - extension
```

Copyright © 2026 Accellera. All rights reserved.

Action `B` defines its own exec declaration `C`, so it shadows the inherited declarations from `A` (including the extension). The tag has a different value than `A`; therefore, only its block is generated:

```
// Action B
```

Action `C` inherits both execs from `A` (original and extension), which share the same tag (`"A"`). These match with the execs already produced by action `A`. Since identical-tag blocks across instances shall be identical, one instance contributes while the others are suppressed. Therefore, `C`'s contribution is suppressed and does not appear in the final output.

Action `D` defines its own exec declaration `C` but with the same tag (`"A"`) as action `A` but with different content, so this results in an error.

20.6 Target-template implementation for functions

When integrating with languages that do not have the concept of a "function," such as assembly language, the implementation for functions can be provided by target-template code strings.

The target-template form of functions (see Syntax 99) allows interactions with a foreign language that do not involve a procedural interface. Examples are injecting assembly code or global variables into generated tests. The target-template forms of functions are always target implementations. Variable references may only be used in expression positions. Function return values shall not be provided, i.e., only functions that return **void** are supported. If a target-template function is an instance (non-static) function, PSS expressions embedded in the target code (using mustache notation) may make reference to the instance attributes, optionally using **`this`**. PSS comments can be added using the hash-inside-braces notation.

See also 20.5.3 and 4.7.1.

20.6.1 Syntax

```
target_template_function ::= target language_identifier [ static ]
  function function_prototype = string_literal ;
```

Syntax 99—Target-template function implementation

The following also apply:

a) Parameter direction shall be unspecified in the function prototype for target-template functions. This implies that the parameter direction (**input**, **output**, or **inout**) shall not be used. If the function declaration contains directions for parameters, this function shall not have a target-template implementation.

b) The prototype specified in the target template declaration shall match the prototype specified in the **function** declaration in the following way:

   1) The number of parameters shall be identical.
   2) The parameter names and types shall be identical.
   3) The return types shall be identical.

20.6.2 Examples

Example 311 provides an assembly-language target-template code block implementation for the `do_stw` function. Function parameters are referenced using mustache notation (**`{{variable}}`**).

Copyright © 2026 Accellera. All rights reserved.

```
package thread_ops_asm_pkg {
    target ASM function void do_stw(bit[32] val, bit[32] vaddr) = """
      loadi RA {{val}}
      store RA {{vaddr}}
    """;
}
```

Example 311—Target-template function implementation

20.7 Procedural constructs

This section specifies the procedural control flow constructs. When relevant, these constructs have the same syntax and execution semantics as the corresponding activity control flow statements (see 11.4).

20.7.1 Scoped blocks

A scoped block creates a new unnamed nested scope, similar to C-style blocks.

20.7.1.1 Syntax

```
procedural_stmt ::=
    procedural_sequence_block_stmt
  | ...
procedural_sequence_block_stmt ::= [ sequence ] { { procedural_stmt } }
```

Syntax 100—Procedural block statement

The **sequence** keyword before the block statement is optional, and is provided to let users state explicitly that the statements are executed in sequence.

Typically, blocks are used to group multiple statements that are part of a control flow statement (such as **repeat**, **if-else**, etc.). It is also valid to have a stand-alone block that is not part of a control flow statement, in which case the following equivalencies apply:

— A stand-alone block that does not create new variables (and hence does not destroy any variables when the scope ends) is equivalent (as far as the constructed AST is concerned) to the case where the contents of the code block are merged with the enclosing parent block. For example:

```
{
    int a;
    int b;
    {
        b = a;
    }
}
```

is equivalent to

```
{
    int a;
    int b;
    b = a;
}
```

— If the start of an enclosing block coincides with the start of the stand-alone nested block (i.e., with no statements in between) and similarly the end of that enclosing block coincides with the end of the

Copyright © 2026 Accellera. All rights reserved.

stand-alone nested block, it is then equivalent to the case where there is just a single code-block with the contents of the nested block. For example:

```
{
    {
        int a;
        int b;
        //
    }
}
```

is equivalent to

```
{
    int a;
    int b;
    //
}
```

20.7.2 Variable declarations

Variables may be declared with the same notation used in other declarative constructs (e.g., **action**). The declaration may be placed at any point in a scope (i.e., C++ style) and does not necessarily have to be declared at the beginning of a scope. However, the declaration shall precede any reference to the variable.

All data types listed in Clause 7 may be used for variable types. It shall be an error to instantiate **rand** variables in a procedural context.

20.7.2.1 Syntax

```
procedural_stmt ::=
    procedural_sequence_block_stmt
  | procedural_data_declaration
  | ...
procedural_data_declaration ::= data_type procedural_data_instantiation
    { , procedural_data_instantiation } ;
procedural_data_instantiation ::= identifier { array_dim } [ = expression ]
```

Syntax 101—Procedural variable declaration

20.7.3 Assignments

Assignments to variables in the scope may be made.

Copyright © 2026 Accellera. All rights reserved.

20.7.3.1 Syntax

```
procedural_stmt ::=
    procedural_sequence_block_stmt
  | procedural_data_declaration
  | procedural_assignment_stmt
  | ...
procedural_assignment_stmt ::= ref_path assign_op expression ;
```

Syntax 102—Procedural assignment statement

The following rules apply to assignments in native PSS functions and execs:

a) A plain-data variable declared within a function/exec scope may be assigned in the scope where it is visible with no restriction.

b) A native PSS function definition may set data attributes of **component** instances through **component** references passed as parameters. Instance functions may similarly set data attributes of their context **component** directly. Since **component** attributes can only be set during the initialization phase, a function that sets such data attributes shall be called only from within **exec init_down** or **init_up**.

c) An **exec init_down** or **init_up** block may set the data attributes of the **component** instance directly in the body of the **exec**.

d) Data attributes of a **struct** instance may be set using the handle passed as a parameter. Similarly, data attributes of **actions** and flow/resource objects may be set using the reference passed as a parameter. A function that sets such data attributes may be invoked in **init**, **solve** or **body** execs.

e) A **struct** instance may be assigned to another **struct** instance of the same type, which results in a deep-copy operation of the data attributes. That is, this single assignment is equivalent to individually setting data attributes of the left-side instance to the corresponding right-side instance, for all the data attributes directly present in that type or in a contained **struct** type. A **struct** instance may be assigned from another **struct** instance that is of a type that inherits from the type of the left-hand side of the assignment. This results in a deep copy of all data attributes present in the base **struct** type (left-hand type) from the right-hand **struct** instance to the left-hand **struct** instance. See 8.3.

20.7.4 Void function calls

Functions not returning a value (declared with **void** return type) may only be called as standalone procedural statements. Functions returning a value may be used as a standalone statement and the return value discarded by casting the function call to **void**:

```
(void)function_call();
```

Calling a nonvoid function as if has no return value shall be legal, but it is recommended to explicitly discard the return value by casting the function call to **void**, as shown above.

Copyright © 2026 Accellera. All rights reserved.

20.7.4.1 Syntax

```
procedural_stmt ::=
    procedural_sequence_block_stmt
  | procedural_data_declaration
  | procedural_assignment_stmt
  | procedural_void_function_call_stmt
  | ...
procedural_void_function_call_stmt ::= [ ( void ) ] function_call ;
```

Syntax 103—Void function call

20.7.5 Return statement

PSS functions shall return a value to the caller using the **return** statement. In PSS functions that do not return a value, the **return** statement without an argument shall be used.

The **return** statement without an argument can also be used in **exec**s. The **return** signifies end of execution—no further statements in the **exec** are executed.

20.7.5.1 Syntax

```
procedural_stmt ::=
    procedural_sequence_block_stmt
  | procedural_data_declaration
  | procedural_assignment_stmt
  | procedural_void_function_call_stmt
  | procedural_return_stmt
  | ...
procedural_return_stmt ::= return [ expression ] ;
```

Syntax 104—Procedural return statement

20.7.5.2 Examples

```
target function int add(int a, int b) {
  return (a+b);
}
```

Example 312—Procedural return statement

20.7.6 Repeat (count) statement

The procedural **repeat** statement allows the specification of a loop consisting of one or more procedural statements. This section describes the *count-expression* variant (see Syntax 105) and 20.7.7 describes the *while-expression* variants.

Copyright © 2026 Accellera. All rights reserved.

20.7.6.1 Syntax

```
procedural_stmt ::=
    procedural_sequence_block_stmt
  | procedural_data_declaration
  | procedural_assignment_stmt
  | procedural_void_function_call_stmt
  | procedural_return_stmt
  | procedural_repeat_stmt
  | ...
procedural_repeat_stmt ::=
    repeat ( [ index_identifier : ] expression ) procedural_stmt
  | ...
```

Syntax 105—Procedural repeat-count statement

The following also apply:

a) *expression* shall be a non-negative integer expression (**int** or **bit**).

b) Intuitively, the *procedural_stmt* is iterated the number of times specified in the *expression*. An optional index-variable identifier can be specified that ranges between `0` and one less than the iteration count. If the expression evaluates to `0`, the *procedural_stmt* is not evaluated at all.

c) The index variable is read-only. Its value shall not be changed within the loop body.

20.7.6.2 Example

```
target function int sum(int a, int b) {
  int res;

  res = 0;

  repeat(b) {
     res = res + a;
  }

  return res;
}
```

Example 313—Procedural repeat-count statement

20.7.7 Repeat-while statement

The procedural **repeat** statement allows the specification of a loop consisting of one or more procedural statements. This section describes the *while-expression* variants (see Syntax 106).

Copyright © 2026 Accellera. All rights reserved.

20.7.7.1 Syntax

```
procedural_stmt ::=
    procedural_sequence_block_stmt
  | procedural_data_declaration
  | procedural_assignment_stmt
  | procedural_void_function_call_stmt
  | procedural_return_stmt
  | procedural_repeat_stmt
  | ...
procedural_repeat_stmt ::=
  . . .
  | repeat procedural_stmt while ( expression ) ;
  | while ( expression ) procedural_stmt
```

Syntax 106—Procedural repeat-while statement

The following also apply:

a) *expression* shall be of type **bool**.

b) Intuitively, the *procedural_stmt* is iterated so long as the *expression* condition is *true*, as sampled before the *procedural_stmt* (in the **while** variant) or after (in the **repeat-while** variant).

20.7.7.2 Example

```
target function bool get_parity(int n) {
  bool parity;

  parity = false;
  while (n != 0) {
     parity = !parity;
     n = n & (n-1);
  }

  return parity;
}
```

Example 314—Procedural while statement

20.7.8 Foreach statement

The procedural **foreach** statement allows the specification of a loop that iterates over the elements of a collection (see Syntax 107).

Copyright © 2026 Accellera. All rights reserved.

20.7.8.1 Syntax

```
procedural_stmt ::=
    procedural_sequence_block_stmt
  | procedural_data_declaration
  | procedural_assignment_stmt
  | procedural_void_function_call_stmt
  | procedural_return_stmt
  | procedural_repeat_stmt
  | procedural_foreach_stmt
  | ...
procedural_foreach_stmt ::=
  foreach ( [ iterator_identifier : ] expression  [ [ index_identifier ] ] ) procedural_stmt
```

Syntax 107—Procedural foreach statement

The following also apply:

a) *expression* shall be of a collection type (i.e., **array**, **list**, **map** or **set**). *expression* may also be an array of *modeling elements* (not including monitors).

b) The body of the **foreach** statement is a sequential block in which *procedural_stmt* is evaluated once for each element in the collection.

c) *iterator_identifier* specifies the name of an iterator variable of the collection element type. Within *procedural_stmt*, the iterator variable, when specified, is an alias to the collection element of the current iteration.

d) *index_identifier* specifies the name of an index variable. Within *procedural_stmt,* the index variable, when specified, corresponds to the element index of the current iteration.

   1) For **array**s and **list**s, the index variable shall be a variable of type **int**, ranging from `0` to one less than the size of the collection variable, in that order.

   2) For **map**s, the index variable shall be a variable of the same type as the **map** keys, and range over the values of the keys. The order of key traversal is undetermined.

   3) For **set**s, an index variable shall not be specified.

e) Both the index and iterator variables, if specified, are implicitly declared within the **foreach** scope and limited to that scope. Regular name resolution rules apply when the implicitly declared variables are used within the **foreach** body. For example, if there is a variable in an outer scope with the same name as the index variable, that variable is shadowed (masked) by the index variable within the **foreach** body. The index and iterator variables are not visible outside the **foreach** scope.

f) Either an index variable or an iterator variable or both shall be specified. For a **set**, an iterator variable shall be specified, but not an index variable.

g) The index and iterator variables are read-only. Their values shall not be changed within the **foreach** body. It shall be an error to change the contents of the iterated collection variable with the **foreach** body.

20.7.9 If-else statement

The procedural **if-else** statement introduces a branch point (see Syntax 108).

Copyright © 2026 Accellera. All rights reserved.

20.7.9.1 Syntax

```
procedural_stmt ::=
    procedural_sequence_block_stmt
  | procedural_data_declaration
  | procedural_assignment_stmt
  | procedural_void_function_call_stmt
  | procedural_return_stmt
  | procedural_repeat_stmt
  | procedural_foreach_stmt
  | procedural_if_else_stmt
  | ...
procedural_if_else_stmt ::= if ( expression ) procedural_stmt [ else procedural_stmt ]
```

Syntax 108—Procedural if-else statement

expression shall be of type **bool**.

20.7.9.2 Example

```
target function int max(int a, int b) {
  int c;

  if (a > b) {
     c = a;
  } else {
     c = b;
  }

  return c;
}
```

Example 315—Procedural if-else statement

20.7.10 Match statement

The procedural **match** statement specifies a multi-way decision point that tests whether an expression matches one of a number of other expressions and executes the matching branch accordingly (see Syntax 109).

Copyright © 2026 Accellera. All rights reserved.

20.7.10.1 Syntax

```
procedural_stmt ::=
    procedural_sequence_block_stmt
  | procedural_data_declaration
  | procedural_assignment_stmt
  | procedural_void_function_call_stmt
  | procedural_return_stmt
  | procedural_repeat_stmt
  | procedural_foreach_stmt
  | procedural_if_else_stmt
  | procedural_match_stmt
  | ...
procedural_match_stmt ::=
   match ( match_expression ) { procedural_match_choice { procedural_match_choice } }
match_expression ::= expression
procedural_match_choice ::=
    [ open_range_list ] : procedural_stmt
  | default : procedural_stmt
```

Syntax 109—Procedural match statement

The following also apply:

a) When the **match** statement is evaluated, the *match_expression* is evaluated.

b) After the *match_expression* is evaluated, the *open_range_list* of each *procedural_match_choice* shall be compared to the *match_expression*. *open_range_list*s are described in 8.5.9.1.

c) If there is exactly one match, then the corresponding branch shall be evaluated.

d) It shall be an error if more than one match is found for the *match_expression*.

e) If there are no matches, then the **default** branch, if provided, shall be evaluated.

f) The **default** branch is optional. There may be at most one **default** branch in the **match** statement.

g) If a **default** branch is not provided and there are no matches, it shall be an error.

20.7.10.2 Example

```
target function int bucketize(int a) {
  int res;

  match (a) {
     [0..3]:  res = 1;
     [4..7]:  res = 2;
     [8..15]: res = 3;
     default: res = 4;
  }

  return res;
}
```

Example 316—Procedural match statement

Copyright © 2026 Accellera. All rights reserved.

20.7.11 Break/continue statement

The procedural **break** and **continue** statements allow for additional control in loop termination (see Syntax 110).

20.7.11.1 Syntax

```
procedural_stmt ::=
    procedural_sequence_block_stmt
  | procedural_data_declaration
  | procedural_assignment_stmt
  | procedural_void_function_call_stmt
  | procedural_return_stmt
  | procedural_repeat_stmt
  | procedural_foreach_stmt
  | procedural_if_else_stmt
  | procedural_match_stmt
  | procedural_break_stmt
  | procedural_continue_stmt
  | ...
procedural_break_stmt ::= break ;
procedural_continue_stmt ::= continue ;
```

Syntax 110—Procedural break/continue statement

The following also apply:

a) The semantics are similar to **break** and **continue** in C++.

b) **break** and **continue** may only appear within loop statements (**repeat**-count, **repeat-while** or **foreach**). Within a loop, **break** and **continue** may be nested in conditional branch or **match** statements.

c) **break** and **continue** affect the innermost loop statement they are nested within.

d) **break** signifies that execution should continue from the statement after the enclosing loop construct. **continue** signifies that execution should proceed to the next loop iteration.

Copyright © 2026 Accellera. All rights reserved.

20.7.11.2 Example

```
// Sum all elements of 'a' that are even, starting from a[0], except those
// that are equal to 42. Stop summation if the value of an element is 0.

function int sum(array<int,100> a) {
   int res;

   res = 0;

   foreach (el : a) {
      if (el == 0)
         break;
      if (el == 42)
         continue;
      if ((el % 2) == 0) {
         res = res + el;
      }
   }

   return res;
}
```

Example 317—Procedural foreach statement with break/continue

20.7.12 Randomize statement

The procedural **randomize** statement shall randomize the specified data attributes or variables.

20.7.12.1 Syntax

```
procedural_stmt ::=
    procedural_sequence_block_stmt
  | procedural_data_declaration
  | procedural_assignment_stmt
  | ...
  | procedural_randomization_stmt
  | procedural_compile_if
  | annotation_usage procedural_stmt
  | stmt_terminator
procedural_randomization_stmt ::=
    randomize procedural_randomization_target procedural_randomization_term
procedural_randomization_target ::= hierarchical_id { , hierarchical_id }
procedural_randomization_term ::=
    with constraint_set
  | ;
```

Syntax 111—Procedural randomize statement

The rules and semantics of the **randomize** statement are described in 13.4.6.

Copyright © 2026 Accellera. All rights reserved.

20.7.13 Yield statement

The target exec blocks of all actions assigned to a single executor execute in a cooperative manner. The **yield** statement temporarily suspends the currently-running exec block, allowing other exec code running in parallel on the same executor to be executed.

20.7.13.1 Syntax

```
procedural_stmt ::=
    procedural_sequence_block_stmt
  | procedural_data_declaration
  | procedural_assignment_stmt
  | ...
  | procedural_yield_stmt
  | stmt_terminator
procedural_yield_stmt ::=
  yield ;
```

Syntax 112—Procedural yield statement

The following also apply:

a) The **yield** statement may only be used in target exec blocks and functions.

b) If no other exec code is currently being executed in parallel, this statement has no effect.

c) If other exec code is currently being executed in parallel, code in at least one other exec block will be executed before the statement after this one executes.

20.7.14 Exec block

Example 318 shows how an **exec body** can be specified using procedural constructs in PSS.

```
action A {
   rand bool flag;

   exec body {
      int var;

      if(flag) {
         var = 10;
      } else {
         var = 20;
      }
      // send_cmd is an imported function
      send_cmd(var);
   }
}
```

Example 318—Exec block using procedural control flow statements

Copyright © 2026 Accellera. All rights reserved.

20.8 Blocking calls and concurrent execution

A PSS tool is expected to enable concurrent execution of multiple exec blocks concurrently-assigned to the same executor via cooperative or preemptive multitasking (21.7.2). PSS provides the *yield* statement (20.7.13) to allow procedural code in exec blocks to explicitly yield control to other concurrently-executing exec blocks assigned to the same executor.

When target exec code blocks on a channel operation (get, put) (see 21.9.1), other concurrently-executing exec blocks assigned to the same executor shall continue to be evaluated.

In Example 319, two actions, `DoDma` and `DoPoll`, are scheduled in parallel on the same executor. `DoDma` blocks waiting for a channel to yield a data item. `DoPoll` polls a register and explicitly yields control. The exec code inside `DoPoll` continues to be evaluated as long as the loop condition is true, including while the `DoDma` action is waiting for the channel to yield data.

Note that order in which blocking exec blocks are awakened and evaluated is non-deterministic.

```
component dma_c {
  ref reg_c<bit>   done;
  channel_c<int> irq;
  action DoDma {
    exec body {
      // ...setup transfer
      // wait for completion interrupt
      irq.get();
    }
  }
  action DoPoll {
    exec body {
      // Expect this loop to continue running
      while (comp.done.read_val() == 0) {
          yield;
      }
    }
  }
  action DmaXferAndPoll {
    activity {
      parallel {
        do DoDma;
        do DoPoll;
      }
    }
  }
}
```

Example 319—Test-realization concurrency

20.9 Comparison between mapping mechanisms

Previous sections describe three mechanisms for mapping PSS entities to external (non-PSS) definitions: functions that directly map to foreign API (see 20.4), functions that map to foreign language procedural code using target code templates (see 20.6), and *exec block*s where arbitrary target code templates are in-lined (see 20.5). These mechanisms differ in certain respects and are applicable in different flows and situations. This section summarizes their differences.

Copyright © 2026 Accellera. All rights reserved.

PSS tests may need to be realized in different ways in different flows:

- by directly exercising separately-existing environment APIs via procedural linking/binding;
- by generating code once for a given model, corresponding to entity types, and using it to execute scenarios; or
- by generating dedicated target code for a given scenario instance.

Table 28 shows how these relate to the mapping constructs.

Table 28—Flows supported for mapping mechanisms

	No target code generation	**Per-model target code generation**	**Per-test target code generation**	**Non-procedural binding**
Direct-mapped functions	X	X	X	
Target-template functions		X	X	
Target-template exec-blocks			X	X

Not all mapping forms can be used for every **exec** kind. Solving/generation-related code shall have direct procedural binding since it is executed prior to possible code generation. *exec block*s that expand declarations and auxiliary files shall be specified as target-templates since they expand non-procedural code. The **run_start** *exec block* is procedural in nature, but involves up-front commitment to the behavior that is expected to run.

Table 29 summarizes these rules.

Table 29—Exec block kinds supported for mapping mechanisms

	Action runtime behavior exec blocks: body	**Non-procedural exec blocks: header, declaration, file**	**Global test exec blocks: run_start, run_end**	**Solve exec blocks: init_down, init_up, pre_solve, post_solve, pre_body**
Direct-mapped functions	X		X (only in pre-generation)	X
Target-template functions	X		X (only in pre-generation)	
Target-template exec-blocks	X	X	X	

The possible use of **action** and **struct** attributes differs between mapping constructs. Explicitly declared prototypes of **function**s enable the type-aware exchange of values of all data types. On the other hand, free parameterization of uninterpreted target code provides a way to use attribute values as target-language meta-level parameters, such as types, variables, functions, and even preprocessor constants.

Copyright © 2026 Accellera. All rights reserved.

Table 30 summarizes the parameter passing rules for the different constructs.

Table 30—Data passing supported for mapping mechanisms

	Back assignment to PSS attributes	Passing user-defined and aggregate data types	Using PSS attributes in non-expression positions
Direct-mapped functions	X	X	
Target-template functions		X	
Target-template exec-blocks			X

20.10 Exported actions

Imported functions and classes specify functions and classes external to the PSS description that can be called from the PSS description. Exported actions specify actions that can be called from a foreign language. See also Syntax 113.

20.10.1 Syntax

```
export_action ::= export [ platform_qualifier ] action_type_identifier
    function_parameter_list_prototype ;
```

Syntax 113—Export action declaration

The **export** statement for an **action** specifies the action to export and the parameters of the action to make available to the foreign language, where the parameters of the exported action are associated by name with the action being exported. The **export** statement also optionally specifies in which phases of test generation and execution the exported action will be available.

The following also apply:

a) As with imported functions (see 20.2.1), the exported action is assumed to always be available if the function availability is not specified.

b) Each call into an **export** action infers an independent tree of actions, components, and resources.

c) Constraints and resource allocation are considered within the inferred action tree and are not considered across imported function / exported action call chains.

20.10.2 Example

Example 320 shows an exported action. In this case, the action `comp::A1` is exported. The foreign language invocation of the exported action supplies the value for the `mode` field of action `A1`. The PSS processing tool is responsible for selecting a value for the `val` field. Note that `comp::A1` is exported to the target, indicating the target code can invoke it.

Copyright © 2026 Accellera. All rights reserved.

```
component comp {

  action A1 {
    rand bit          mode;
    rand bit[32]    val;

    constraint {
      if (mode!=0) {
        val in [0..10];
      } else {
        val in [10..100];
      }
    }
  }

}

package pkg {
  // Export A1, providing a mapping to field 'mode'
  export target comp::A1(bit mode);
}
```

Example 320—Export action

20.10.3 Export action foreign language binding

An exported action is exposed as a function in the target foreign language (see Example 321). The component namespace is reflected using a language-specific mechanism: C++ namespaces, SystemVerilog packages. Parameters to the exported action are implemented as parameters to the foreign language function.

```
namespace comp {
   void A1(unsigned char mode);
}
```

Example 321—Export action foreign language implementation

Copyright © 2026 Accellera. All rights reserved.

21. PSS core library

The PSS *core library* provides standard portable functionality and utilities for common PSS applications. It defines a set of **component** types, data types, **functions**, and attributes. The interface of the core library is specified in PSS-language terms, and its use conforms to the rules of the language. However, the full semantics of its entities involve reference to type information, solving, scheduling, and runtime services. Hence, the implementation of the core library depends on inner workings of PSS processing tools and is expected to be coupled with them.

The core library currently covers functionality in the following areas:

- String formatting and output operations
- File operations
- Error reporting
- Randomization
- Manipulation and storage of floating-point values
- Representation of execution contexts in the target environment
- Assignments of actions and flow/resource objects to execution contexts
- Representation of target address spaces
- Allocation from and management of target address spaces
- Access to target address spaces
- Representation of and access to registers
- Synchronization and communication
- Standard annotations

The core library functionality is defined in four packages:

- **`std_pkg`**, covering string formatting, file operations, error reporting, randomization, and core data types
- **`executor_pkg`**, covering representation of execution contexts and assignment of actions and flow/resource objects to execution contexts
- **`addr_reg_pkg`**, covering representation of address spaces and access to memory, and representation and access to registers
- **`sync_pkg`**, covering synchronization and communication constructs

This section covers the interface, semantics, and intended use of core library entities in the areas listed above. Note that it defines a library interface, not new language constructs. The code for the built-in library package contents appears in Annex C.

In the following sections, library code definitions may omit reiterating the surrounding package, and example code may omit importing core library packages for brevity.

21.1 String formatting and output

The PSS core library provides means for string formatting and output operations. The built-in package **`std_pkg`** defines functions and types for these purposes, as well as for file operations and error reporting, introduced in the next two sections of this document.

On solve platforms, a complete set of input/output and file operations is provided, similar to other programming languages. Functions are defined for string formatting, printing, and reading from and writing to text files.

Copyright © 2026 Accellera. All rights reserved.

On target platforms, a limited portable messaging capability is provided, because some target environments may not have a file system or access to string manipulation libraries such as in C.

21.1.1 String formatting

Several output functions involve a string formatting capability. They are based on an approach similar to C printf()-style string formatting. Each of these functions gets a *format string* parameter **`format_str`** of type **string**, followed by a generic varargs parameter **`args`**.

The format string is used as a template, where all characters are taken literally except when the character `%` appears. A `%` followed by another `%` denotes a single literal `%`. Otherwise, a `%` starts a *format specifier*.

A format specifier determines how data passed in each subsequent function parameter (passed as varargs) should be embedded in the resulting string. It consists of the following optional parts followed by a *formatting character*:

```
%[flags][width][.precision]format
```

The optional *flags*, if specified, denote the following:

`-`	Left justification (default is right justification)
`+`	Force a sign (+ or -) to precede numeric values. By default, positive numbers are not preceded with +.
`space`	If a numeric value is not preceded by a sign, it is preceded by a space.
`#`	For **`o`**, **`x`**, **`X`**, **`b`** or **`B`** format characters, the value is preceded with **`0`**, **`0x`**, **`0X`**, **`0b`** or **`0B`**, respectively, for values different from zero. For floating-point formats, force a decimal point even if no more digits follow the decimal point.
`0`	When left padding is used, pad a numeric value with zeros instead of spaces.

The optional *width*, if specified, denotes the minimum number of characters to insert into the formatted string. The inserted value is not truncated if larger than the specified width. *width* is typically used to pad fixed-width fields in tabulated output.

The optional *precision*, if specified, denotes the following:

- For integer formats (including **`p`**), specifies the minimum number of digits to be inserted into the formatted string. If needed, the result is padded with leading zeros. The value is not truncated even if the result is longer. A *precision* of `0` means that no character is inserted for the value `0`.
- For floating-point formats **`e`**, **`E`**, and **`f`**, specifies the number of digits to be inserted *after* the decimal point. By default, this is 6.
- For floating-point formats **`g`** and **`G`**, specifies the maximum number of significant digits to be inserted.
- For **`s`** and **`n`** formats, specifies the maximum number of characters to be inserted. By default, all characters in the string, the enumeration item name, or the boolean name are used. Truncation, if needed, is from the right.

Copyright © 2026 Accellera. All rights reserved.

If *precision* is empty (the period is specified without an explicit value for *precision*), `0` is assumed.

The *formatting character* determines the expected data type of the corresponding function parameter and how it is formatted, as follows:

`d`	A signed integer in decimal radix.
`u`	An unsigned integer in decimal radix.
`x,X`	An unsigned integer in hexadecimal radix. **`x`** uses lowercase letters and **`X`** uses uppercase.
`o`	An unsigned integer in octal radix.
`b,B`	An unsigned integer in binary radix. If **`#`** flag is specified, **`b`** uses lowercase **`0b`** and **`B`** uses uppercase **`0B`**.
`f`	A floating-point value in decimal form. For example, `123.4567`.
`e,E`	A floating-point value in scientific form. For example, `1.234567e+02`. **`e`** uses lowercase `e` for the exponent and **`E`** uses uppercase `E`.
`g,G`	A floating-point value in the shortest form, decimal or scientific. If scientific form, **`g`** corresponds to **`e`**, and **`G`** corresponds to **`E`**.
`n`	An enumeration item value in the form of its name, or a Boolean value in the form of "`false`" or "`true`".
`s`	A string.
`p`	A chandle as a pointer value in hexadecimal form, including the preceding **`0x`** (similar to **`%#x`** for integer numbers)

The following also apply:

a) If the format string contains `%` followed neither by another `%` nor by a valid format specifier, an error shall be generated.

b) The number of format specifiers in the format string shall be equal to the number of parameters in the varargs. Otherwise, an error shall be generated.

c) Each format specifier in the format string shall match the type of the corresponding parameter in the varargs. Implicit type conversions shall be allowed. For example: if **`%d`** is used for a parameter of an unsigned type, the value is converted to signed type before being formatted; if **`%f`** is used for a parameter of an integer type, the value is converted to floating-point before being formatted; if **`%d`** is used for a parameter of a floating-point type, the value is converted to an integer before being formatted. There is one exception to this rule: unsigned integer formats (**`%u`**, **`%x`**, **`%X`**, **`%o`**, **`%b`**, **`%B`**) shall not be allowed for floating-point values because there is no well-defined conversion from a negative floating-point value to a positive integer without a specific width. If the type does not match and an implicit type conversion is not applicable, an error shall be generated.

d) If an enumeration item is declared using an escaped identifier (see 4.3), the leading backslash is not considered as a part of the name. Therefore, its formatting with %n does not contain the leading backslash.

Copyright © 2026 Accellera. All rights reserved.

21.1.2 Solve-time string formatting and output

The functions **`format()`** and **`print()`** are used on the solve platform to facilitate the string formatting functionality. The function **`format()`** returns a formatted string. The function **`print()`** outputs a formatted string to the standard output and can be used to display and log certain information.

```
package std_pkg {
  solve pure function string format(string format_str, type... args);
  solve function void print(string format_str, type... args);
}
```

Syntax 114—String formatting and output functions

Example 322 demonstrates how native functions can be used to print or to return a formatted string of the context of a given struct instance.

```
import std_pkg::*;
struct my_struct {
   int value;
   string name;
}
solve function void print_foo(my_struct s) {
   print("The context of the struct is:\n");
   print("value = %d\nname = '%s'\n", s.value, s.name);
}
solve function string get_foo_context_string(my_struct s) {
   return format("value = %d\nname = '%s'\n", s.value, s.name);
}
```

Example 322—Printing or formatting the context of a struct

21.1.3 Message logging

The function **`message()`** is used to log certain information during the solving and execution of a test in a portable way. It inserts a text line, including a trailing newline ('**`\n`**'), into the execution log.

```
package std_pkg {
  enum message_verbosity_e {NONE, LOW, MEDIUM, HIGH, FULL};
  function void message (message_verbosity_e vrb_level, string format_str, type... args);
}
```

Syntax 115—Runtime messaging function

The PSS processing tool shall provide means for specifying a messaging verbosity level for a given test run. For a higher test run verbosity level, more messages will be issued and more information will be provided.

The parameter **`vrb_level`** denotes the verbosity level of a particular message, and determines the minimum test run verbosity level for which the message should be issued. Messages with verbosity higher than the test run verbosity level shall be ignored.

Copyright © 2026 Accellera. All rights reserved.

For example, a message of verbosity level **NONE** is considered non-verbose; it is typically a critical message which shall always be issued regardless of the verbosity level of the run. A message of verbosity level **LOW** shall not be issued in a run whose verbosity level is **NONE**, but shall be issued in all other cases, because it is typically an important, though not critical, message. A message of verbosity level **FULL** is considered very verbose, and it shall only be issued in a run whose verbosity level is **FULL**; it is typically a least important message which may provide some additional details or information which is not essential in most runs.

In some environments (for example, certain embedded software environments), restrictions may apply to the *vrb_level* and *format_str* expressions when the message function is called in a target exec context. This is due to the target platform memory requirements for the string operations or other considerations.

If expressions with side effects, such as non-pure function calls, are passed as parameters to **message()**, their evaluation is not guaranteed because the verbosity level of a particular test run may determine whether they are evaluated. Therefore, users should not rely on the side effects of such expressions as parameters to **message()**.

Example 323 demonstrates the usage of **message()** in an **exec body** block. There are two messages: the first message of verbosity level **FULL**, and the second message of verbosity level **LOW**. In test runs whose verbosity level is **NONE**, no message is issued. In runs whose verbosity level is at least **LOW** but lower than **FULL**, only the second message is issued. In runs with verbosity level **FULL**, both messages are issued.

```
import std_pkg::*;
component C {
   target function int my_func() {…}
   action A {
      rand int x;
      exec body {
         int y = comp.my_func();
         message(FULL, "The values of the variables x and y are: ");
         message(LOW, "%d, %d", x, y);
      }
   }
}
```

Example 323—Runtime messages

21.2 File operations

The PSS core library provides two flavors of text input/output operations on solve platform files. Files can be opened separately to obtain a file handle, which can then be used when calling write and read functions. Alternatively, write and read can be performed with a single function call that also opens and closes the file.

File read and write operations in both flavors use string values.

Syntax 116 specifies types and functions used for file operations that use file handles.

Copyright © 2026 Accellera. All rights reserved.

```
package std_pkg {
  typedef chandle file_handle_t;
  static const file_handle_t nullfilehandle = /* implementation-specific */;

  enum file_option_e {TRUNCATE, APPEND, READ};

  solve function file_handle_t file_open(string filename, file_option_e opt);
  solve function void file_close(file_handle_t file_handle);
  solve function bool file_exists(string filename);

  solve function void file_write
     (file_handle_t file_handle, string format_str, type... args);
  solve function string file_read(file_handle_t file_handle, int size = -1);
  }
```

Syntax 116—Text file operations using file handles

The type **file_handle_t** is used to represent a text file that is open for the purpose of reading or writing. A file handle is obtained by calling function **file_open()**.

Values of the enumeration type **file_option_e** represent the purpose of the file, as follows:

- **TRUNCATE** Delete any existing content of the file and allow write operations.
- **APPEND** Allow write operations; text will be appended to the existing file content.
- **READ** Allow read operations.

The function **file_open()** returns a file handle to the file whose name is specified by **filename**. If the file fails to open in the mode specified by **opt**, the special value **nullfilehandle** is returned.

The function **file_close()** closes the file represented by **file_handle**, which shall have been previously opened and not closed. Once a file has been closed, the handle can no longer be used for reading or writing.

The function **file_exists()** returns true if a file with the specified filename exists in the file system, otherwise returns false.

The function **file_write()** writes a formatted string to a file represented by **file_handle**, which shall have been opened with the **TRUNCATE** or **APPEND** option. A newline is not added at the end.

The function **file_read()** reads at most **size** number of characters from a file represented by **file_handle**, and returns a string containing those characters. It starts reading the characters from the beginning of the file (if it is the first call after opening the file), or from the first position not read by a previous **file_read()** invocation. If **size** is negative (or not specified), the content of the file is read till the end. The file shall have been opened with the **READ** option.

The functions **file_write()**, **file_read()**, and **file_close()** shall trigger an appropriate error if the operation cannot be performed.

Syntax 117 specifies functions used for file reading and writing in a single function call.

Copyright © 2026 Accellera. All rights reserved.

```
package std_pkg {
  solve function void file_write_lines
    (string filename, list<string> lines, file_option_e opt);

  solve function list<string> file_read_lines(string filename);
}
```

Syntax 117—Simple text file operations

The function **`file_write_lines()`** writes all strings in **`lines`** to the file whose name is specified by **`filename`**. A newline character is inserted at the end of each string. **`opt`** shall be either **`TRUNCATE`** or **`APPEND`**. If **`APPEND`** is used, a newline character is also inserted at the end of the existing file content, unless the last character in it is already a newline character.

The function **`file_read_lines()`** reads the entire file whose name is specified by **`filename`**, and returns a list of strings representing the text in the file. A string is terminated when a newline character in the file is reached. The newline characters themselves are not included in the strings.

Both functions trigger an appropriate error if the operation cannot be performed.

In principle, the string passed as the **`filename`** parameter to functions **`file_open()`**, **`file_exists()`**, **`file_write_lines()`**, and **`file_read_lines()`** can include a directory path. PSS processing tools may provide specific ways of mapping a physical file in the file system to a given **`filename`** string. For example, a tool may use an environment variable to provide one or more search paths for files (similar to a **`PATH`** environment variable used in many operating systems to search for executable files).

Example 324 shows two functions that write the content of a given struct list into a text file in a certain arbitrary format. Both functions achieve the same result, but the first function uses a file handle, and the second function uses **`file_write_lines()`** directly.

Copyright © 2026 Accellera. All rights reserved.

```
import std_pkg::*;
struct my_struct {
   int value;
   string name;
}
solve function void write_my_struct_list_using_file_handle
                        (string file_name, list<my_struct> s_list)
{
   file_handle_t f = file_open(file_name, TRUNCATE);
   foreach (s: s_list) {
      file_write(f, "%d %s\n", s.value, s.name);
   }
   file_close(f);
}
solve function void write_my_struct_list_using_string_list
                        (string file_name, list<my_struct> s_list)
{
   list<string> lines;
   foreach (s: s_list) {
      lines.push_back(format("%d %s", s.value, s.name));
   }
   file_write_lines(file_name, lines, TRUNCATE);
}
```

Example 324—File operations

21.3 Error reporting

The functions **`error()`** and **`fatal()`** are used to report an error during a test and/or to abort the rest of the run in a portable way. They are similarly used for the solving process.

```
package std_pkg {
  function void error(string format_str, type... args);
  function void fatal(int status, string format_str, type... args);
}
```

Syntax 118—Error reporting functions

Both **`error()`** and **`fatal()`** insert the specified formatted text into the solving or execution log, with a trailing newline character.

The parameter **`format_str`** shall be a string expression whose value is known at solve time. If there are strings (as opposed to numbers) among the subsequent **`args`** data parameters, they shall also be known at solve time. In other words, when used in target contexts, the string values of those parameters shall be constant at run time.

The function **`fatal()`** shall terminate the solving or execution flow at the nearest possible point. The value of the parameter **`status`** is returned to the calling environment.

The function **`error()`** may terminate the solving or execution flow, or it may not, depending on tool/session-specific criteria.

Copyright © 2026 Accellera. All rights reserved.

Example 325 demonstrates reporting of a run-time error under a particular condition.

```
component C {
   function int get_some_id();
   action A {
      exec body {
         int id = comp.get_some_id();
         if (id > 1000) {
            error("Id is too large: %d", id);
         }
      }
   }
}
```

Example 325—Error reporting

21.4 Randomization

Randomization functions are contained within the **std_pkg** package.

```
package std_pkg {
  function bit[32] urandom();
  function bit[32] urandom_range(bit[32] min, bit[32] max);
}
```

Syntax 119—Randomization functions

21.4.1 Urandom()

The **urandom()** function returns an unsigned 32-bit integer.

21.4.2 Urandom_range(min, max)

The **urandom_range()** function returns an unsigned 32-bit integer between the specified minimum and maximum values.

21.5 Floating-point

PSS defines a set of functions for manipulating floating-point values and representing various storage formats of floating-point numbers. These functions and data types are defined in the **std_pkg** package.

Copyright © 2026 Accellera. All rights reserved.

21.5.1 Floating-point storage types

```
struct float_base_s <int Wm, int We, endianness_e E=LITTLE_ENDIAN> :
  packed_s<E> {
    rand bit[Wm] mantissa;
    rand bit[We] exponent;
    rand bit     sign;
  }
typedef float_base_s<23, 8> float32_s;
typedef float_base_s<52,11> float64_s;
```

Syntax 120—Floating-point storage types

The PSS core library defines a **struct**, **`float_base_s`**, to represent the in-memory layout of floating-point numbers. Specific specializations of this templated type are used to capture specific storage layouts. The **`float_base_s`** struct inherits from the **`packed_s`** struct type, which is described in 21.13. Storage formats for the two built-in computation types are defined as part of the core library package.

21.5.2 Floating-point computation functions

The PSS core library defines the following floating-point computation functions. All functions return **float64** as a result, and accept parameters of type **float64**. Their behavior shall match the equivalent C language standard math library function with the same name, since **float64** is equivalent to the **`double`** type in C. Function prototypes may be found in Annex C.

Floating-point functions may not be used in constraints.

Table 31—Floating-point computation functions

Function	Description
`log(x)`	Natural logarithm
`log10(x)`	Decimal logarithm
`exp(x)`	Exponential
`sqrt(x)`	Square root
`pow(x,y)`	x^y
`round(x)`	Round to nearest value
`floor(x)`	Floor
`ceil(x)`	Ceiling
`sin(x)`	Sine
`cos(x)`	Cosine
`tan(x)`	Tangent
`asin(x)`	Arc-sine

Copyright © 2026 Accellera. All rights reserved.

Table 31—Floating-point computation functions (Continued)

Function	Description
`acos(x)`	Arc-cosine
`atan(x)`	Arc-tangent
`atan2(y,x)`	Arc-tangent of y/x
`hypot(x,y)`	sqrt(x*x+y*y)
`sinh(x)`	Hyperbolic sine
`cosh(x)`	Hyperbolic cosine
`tanh(x)`	Hyperbolic tangent
`asinh(x)`	Arc-hyperbolic sine
`acosh(x)`	Arc-hyperbolic cosine
`atanh(x)`	Arc-hyperbolic tangent

21.5.3 Computation-type field extraction and composition

Floating-point computation and storage data types both have a sign, exponent, and mantissa component. Floating-point types differ in the width of the exponent and mantissa components. PSS defines functions for accessing the various components of computation types, and functions for forming a computation-type value from floating-point component parts.

```
pure function bit[52] float_mantissa(float64 fv);
```

Syntax 121—Float_mantissa function

The **`float_mantissa()`** function extracts the mantissa bit image from the specified **float64** value as is with no conversion.

```
pure function bit[11] float_exponent(float64 fv);
```

Syntax 122—Float_exponent function

The **`float_exponent()`** function extracts the exponent bit image from the specified **float64** value as is with no conversion.

```
pure function bit float_sign(float64 fv);
```

Syntax 123—Float_sign function

The **`float_sign()`** function extracts the sign bit of the specified **float64** value.

Copyright © 2026 Accellera. All rights reserved.

```
pure function float64 to_float(bit[52] mantissa, bit[11] exp, bit sign);
```

Syntax 124—To_float function

The **to_float()** function composes a **float64** value from the specified sign, exponent, and mantissa component bit images.

```
typedef float_base_s<7,8> bfloat16_s;

struct S {
   exec post_solve {
     float64 f1 = 20.25;
     bfloat16_s f2;
     float64 f3;
     f2.sign = float_sign(f1);

     // Unbias from 11-bit exponent, and bias for 8-bit
     f2.exponent = float_exponent(f1) - (2**(11-1)-1) + (2**(8-1)-1);

     // Use the leftmost bits, so we lose some precision
     // but preserve the correct value.
     f2.mantissa = float_mantissa(f1) >> (52-7);

     f3 = to_float(
        f2.mantissa << (52-7),
        f2.exponent - (2**(8-1)-1) + (2**(11-1)-1),
        f2.sign);
   }
}
```

Example 326—Conversion to and from storage type

Example 326 above shows conversion of the floating-point value `20.25` held in a **float64** variable to a `bfloat16_s` floating-point storage data type. The `bfloat16_s` storage type has an exponent of 8 bits and a mantissa of 7 bits, while the **float64** variable has an exponent of 11 bits and a mantissa of 52 bits.

In this example, the exponent part is stored in the storage type in biased form. To achieve this, it is first unbiased from the original bit image representation of 11 bits (by subtracting 2^{11-1}-1) and then biased for 8 bits (by adding 2^{8-1}-1). For the mantissa part, the 7 *left-most* bits are used, which is achieved by left-shifting by 52-7 bits.

Finally, the components of the `bfloat16_s` type are converted back to a **float64** value using the **to_float()** function.

21.6 Standard annotations

21.6.1 Element code documentation

The `code_doc` annotation associates implementation-level documentation text with an element of the PSS model. The `code_doc` annotation may be applied to statements and other executable elements of the PSS model.

Copyright © 2026 Accellera. All rights reserved.

PSS processing tools may use the `code_doc` annotation to influence generated implementation artifacts. For example, a tool may emit the associated text as a comment in generated target code. The interpretation and use of this annotation is tool-specific.

```
package std_pkg {
  annotation code_doc {
    string text;
  }
```

Syntax 125—Code_doc annotation

Example 327 shows how the `code_doc` annotation is applied to a statement and Example 328 provides an example of how a tool may choose to use that content.

```
function void zero_mem32(addr_handle_t base, int count) {
  repeat (i : count) {
    @code_doc {.text="Zero the target memory word" }
    write32(make_handle_from_handle(base, 4*i), 0);
  }
}
```

Example 327—Applying the code_doc annotation in PSS

```
void zero_mem32(void *base, int count) {
  int i;
  for (i=0; i<count; i++) {
    // Zero the target memory word
    *(((volatile unsigned int *)base)+4*i) = 0;
  }
```

Example 328—Example of tool output from the code_doc annotation

21.6.2 Element documentation

The `doc` annotation associates model-level documentation text with an element of the PSS model. PSS processing tools may extract and use these annotations.

```
package std_pkg {
  annotation doc {
    string text;
  }
}
```

Syntax 126—Element documentation

21.7 Executors

A PSS generated test calls foreign functions available in the target environment, executes target-language code blocks, and performs target operations provided in the core-library. It does so in accordance with the

Copyright © 2026 Accellera. All rights reserved.

user-defined realization of actions and of flow/resource objects specified in the form of target **exec** blocks—**body**, **run_start**, and **run_end**—and functions called from them. Foreign function calls, target-language code blocks, and built-in target operations, all need to be performed under a certain agent of execution available to the test in the runtime environment, or in short, an *executor*.

An *executor* is an abstract notion that may correspond to different kinds of entities in different environments. For example:

- An embedded processor core or HW thread in a bare-metal environment that executes code generated by the PSS tool
- A BFM instantiated as a master on an interconnect of the DUT that exposes transactional APIs to the PSS tool
- A transactor, or testbench agent, connected to an I/O interface of the system that exposes transactional APIs, or higher-level stimulus sequences, to the PSS tool

The PSS core library provides means to represent executors in the PSS description and to assign scenario entities to them. Executors are characterized by user-defined properties called *traits*, which serve to control the assignment of actions/objects to them. For example, the cluster of a CPU core could be represented as a trait attribute. Related executors are grouped together so that scenario entities can be assigned to a random instance out of a group. The selection of executors satisfies constraints on their trait attributes, if any are specified.

In addition, executors can be used to customize the implementation of target functions for specific environments. Actions assigned to different executors can thereby employ different mappings of portable operations.

The PSS built-in package **`executor_pkg`** defines types and functions related to the management of executors. In subsequent sections, except Syntax 127, the enclosing **`executor_pkg`** is omitted for brevity. Examples may also omit import of **`executor_pkg`**.

21.7.1 Executor representation

An executor is an execution agent or context available to the test in the runtime environment. Executors are represented using a core-library component type instantiated in the PSS description. Actions and flow/resource objects may subsequently be assigned to these executors. This assignment is controlled through an *executor claim struct* (see 21.7.2).

Representing executors in a PSS description is optional. In the absence of executor instances, PSS tools are free to determine the execution context of entities based on other considerations, such as global defaults or policies.

21.7.1.1 Executor component type

An executor is represented using the template component **`executor_c`**, or a subtype of it. The template parameter is used to tag the executor and possibly to provide additional selection attributes. Template **`executor_c`** is derived from **`executor_base_c`**.

Copyright © 2026 Accellera. All rights reserved.

```
package executor_pkg {
  struct executor_trait_s {};

  struct empty_executor_trait_s : executor_trait_s {};

  component executor_base_c {
    target function chandle get_context();
  };

  component executor_c
    <struct TRAIT : executor_trait_s = empty_executor_trait_s>
    : executor_base_c {
     TRAIT trait;
  };
  ...
}
```

Syntax 127—Executor component

An executor component is strictly a test-realization artifact. It shall be an error to declare in its scope scenario model elements, namely: **action** types, **pool** instances, and pool binding directives.

21.7.1.1.1 Get_context function

The `get_context` function returns a handle that is used to identify the executor context when calling a PSS export function.

`get_context` is implemented by the PSS tool. The user may not provide a customized implementation of this method.

21.7.1.2 Executor group component type

Component **`executor_group_c`** is used to group one or more executors that serve similar purposes. Actions and flow/resource objects that *claim* an executor are assigned to an executor selected out of one specific group (see more on matching rules in 21.7.2.2).

```
component executor_group_c
                <struct TRAIT : executor_trait_s = empty_executor_trait_s> {
                  solve function void add_executor(ref executor_c<TRAIT> exe);
};
```

Syntax 128—Executor group component

An executor group component is strictly a test-realization artifact. It shall be an error to declare in its scope scenario model elements, namely: **action** types, **pool** instances, and pool binding directives.

Copyright © 2026 Accellera. All rights reserved.

21.7.1.2.1 Add_executor function

Instance function **add_executor** (see Syntax 128) of **executor_group_c** is used to populate the group with executor instances. Executors added to a group shall all match with the group's trait struct type. The **add_executor** function may only be called in **exec init_down** and **init_up** blocks.

The following also apply:

a) Any executor can be added to a given group, regardless of where it is instantiated in the component instance tree. This includes executors instantiated above the group, below it, or in a different sub-tree.

b) An executor instance may not be added more than once to the same group.

c) An executor instance may be added to more than one group.

d) An executor does not have to be added to any group. An executor that is not part of any group would be inactive—no **exec** blocks would ever be assigned to it.

Example 329 demonstrates how executors are defined, instantiated, and added to an executor group. The executor group `my_hybrid_group_c` is populated with two different executor types. These two types may vary in properties, but are both derived from the instantiation of template **executor_c** with the struct type `master_trait_s`. The executors in this group are treated symmetrically when assigning actions to them.

```
struct master_trait_s : executor_trait_s {};

component my_core_executor_c : executor_c<master_trait_s> { ... };

component my_bus_vip_executor_c : executor_c<master_trait_s> { ... };

component my_hybrid_group_c : executor_group_c<master_trait_s> {
   my_core_executor_c cores[4];
   my_bus_vip_executor_c bfms[2];

   exec init_down {
      foreach (c: cores) {
         add_executor(c);
      }
      foreach (b: bfms) {
         add_executor(b);
      }
   }
};
```

Example 329—Defining an executor group

21.7.2 Executor assignment

An action or a flow/resource object can declare its claim for an executor by instantiating a *claim struct*. Each claim instance is statically matched to an executor group that is nearest in the component instance tree and parameterized by the same trait struct type. The entity is assigned to an executor out of the matching group, which satisfies the trait constraints.

It is not required that scenario entities be explicitly assigned to an executor even if they contain target **exec** blocks. In the absence of explicit assignments, the executor assigned with the containing component of the action will be used (see 21.7.2.6). In the absence of explicit assignment of a component executor, PSS tools

Copyright © 2026 Accellera. All rights reserved.

are free to determine the execution context of entities based on other considerations, such as global defaults or policies.

In the absence of explicit assignments for flow/resource objects the default executor of the component in which the pool is declared will be assigned.

The executor of the containing entity will serve as the executor for struct types.

Executors do not generally limit concurrency of PSS behaviors in a test scenario. In cases where concurrently scheduled actions are assigned to the same underlying executor, the PSS tool is responsible for employing the means to enable concurrent execution, such as preemptive or cooperative multitasking.

21.7.2.1 Executor claim struct type

An action or a flow/resource object can control its assignment to an executor by declaring an *executor claim* —an attribute of template struct type **`executor_claim_s`**. An executor claim can be a direct field of the entity, a field of any of its nested structs, or in the case of flow/resource objects, the supertype from which the object is derived. In all these cases, the assignment to an executor applies in the same way.

An action or a flow/resource object may be assigned to no more than one executor. Therefore, there can only be one executor claim struct anywhere under a given action or object. Multiple executor claim structs within the same action or object shall be flagged as an error. Note that the assignment of executors per an executor claim is not exclusive, and is generally unrelated to the relative scheduling of actions.

```
struct executor_claim_s
        <struct TRAIT : executor_trait_s = empty_executor_trait_s> {
  rand TRAIT trait;
};
```

Syntax 129—Executor claim struct

Example 330 demonstrates the use of the **`executor_claim_s`** struct. In this case, **action** `A` declares an executor claim. `A`'s executor claim is matched with executor group `eg` that is instantiated directly under its context **component** `C`, as both are parameterized with the same (default) trait type. Consequently, **action** `A` is necessarily assigned to the executor `e` instantiated under its context component. **Component** `C` is instantiated twice under `pss_top`. Under the entry **action** `test`, **action** `A` is invoked three times. The generated test will call the function `do_something()` twice under the execution context associated with executor `c1.e`, and subsequently once under the execution context associated with executor `c2.e`.

Copyright © 2026 Accellera. All rights reserved.

```
component C {
   executor_c<> e;
   executor_group_c<> eg;
   exec init_down {
      eg.add_executor(e);
   }

   action A {
      rand executor_claim_s<> ec;
      exec body C = """
         do_something();
      """;
   };
};

component pss_top {
   C c1,c2;

   action test {
      C::A a1, a2, a3;
      activity {
         parallel {
            a1 with { comp == this.comp.c1; };
            a2 with { comp == this.comp.c1; };
            a3 with { comp == this.comp.c2; };
         }
      }
   };
};
```

Example 330—Simple executor assignment

21.7.2.2 Rules for matching an executor claim with an executor group

An executor claim is matched with an executor group for the purpose of selecting an executor. The matching is based on the static structure of the model. A claim is resolved to an executor group that:

a) is parameterized by the same trait type as the claim;

b) is instantiated in a containing component of the declaring scenario entity (the context component hierarchy of an action or the container component of a flow/resource object pool);

c) and is nearest in the component hierarchy going up from the context component to the root component.

It shall be an error if no executor group matches a claim per the above rules. Similarly, it shall be an error if more than one executor group in the component context identified in b) matches a claim.

Note that given the above rules, instantiating a group within a group would be pointless, as no executor claim could match the inner group.

Copyright © 2026 Accellera. All rights reserved.

21.7.2.3 Claim trait semantics

The trait type of an executor claim shall be the same as that of the executor selected for the declaring entity. In addition, the trait attribute values of the executor claim instance shall be equal to the values of the corresponding attributes of the executor trait. Hence, the selected executor shall satisfy the claim trait constraints.

Example 331 demonstrates the use of the executor trait struct for the selection of executors. In this example, executors in group `my_embedded_cores_group_c`, representing eight CPU cores, are classified into two clusters, each consisting of four cores. Action `my_ip_c::op` claims an executor. It constrains the selection of the executor, relating the executor cluster ID to other attributes. Action `ops_on_two_clusters` executes two `op` actions, one on each cluster. Note that the one assigned to cluster `0` will have its input buffer `mem_kind` not equal to `DDR`, due to the constraint in action `op`.

```
struct my_core_trait_s : executor_trait_s {
   rand int in [0..1] cluster_id;
};

component my_embedded_cores_group_c : executor_group_c<my_core_trait_s> {
   executor_c<my_core_trait_s> cores[8];
   exec init_down {
      foreach (c: cores[i]) {
         c.trait.cluster_id = i/4;
         add_executor(c);
      }
   }
};

component my_ip_c {
   action op {
      input data_buff in_buff;
      rand executor_claim_s<my_core_trait_s> core;
      constraint in_buff.mem_kind == DDR -> core.trait.cluster_id != 0;
   };
};

component pss_top {
   my_embedded_cores_group_c embedded_core_group;
   my_ip_c my_ip;

   action ops_on_two_clusters {
      activity {
         do my_ip_c::op with { core.trait.cluster_id == 0; };
         do my_ip_c::op with { core.trait.cluster_id == 1; };
      }
   };
};
```

Example 331—Definition and use of executor trait

Copyright © 2026 Accellera. All rights reserved.

21.7.2.4 Executor resources

In some cases, the assignment of certain actions to executors needs to be exclusive, ruling out the handling of concurrent actions by the same execution agent. Resource claims and resource pools express such rules at the scenario model level, guaranteeing that random schedules satisfy the resource consistency of executors. In these cases, the executor assigned to actions needs to be in strict correspondence with the resource instance claimed by them.

A resource object that is derived from template struct **`executor_claim_s`** is considered a claim not just for the purpose of its own executor assignment, but also for that of the actions that claim it as a resource in either **lock** or **share** mode. In other words, from the executor assignment point of view, a reference to a resource object derived from struct **`executor_claim_s`** functions like an executor claim of the action itself.

In Example 332, resource object `my_core_r` represents a processor core at the scenario model level. Action `my_ip_c::op1` needs to be assigned a core exclusively for its duration, and therefore **lock**s a resource instance. Action `my_ip_c::op2` does not require exclusive use of a core, and therefore claims a resource instance in **share** mode. Action **`test`** executes a random selection of `op1` and `op2`, which need to be scheduled consistently across the different cores.

Copyright © 2026 Accellera. All rights reserved.

```
struct my_core_trait_s : executor_trait_s {
   rand int in [0..7] core_id;
};

resource my_core_r : executor_claim_s<my_core_trait_s> {
   constraint trait.core_id == instance_id;
};

component my_cores_group_c : executor_group_c<my_core_trait_s> {
   executor_c<my_core_trait_s> cores[8];
   exec init_down {
      foreach (c: cores[i]) {
         c.trait.core_id = i;
         add_executor(c);
      }
   }
};

component my_ip_c {
   action op1 {
      lock my_core_r core;
      exec body {
         my_ip_blocking_op();
      }
   };

   action op2 {
      share my_core_r core;
      exec body {
         while (!my_ip_op2_done()) { yield; }
      }
   };
};

component pss_top {
   my_cores_group_c core_group;
   pool [8] my_core_r core_pool;
   bind core_pool *;

   my_ip_c my_ip;

   action test {
      activity {
         schedule {
            replicate (10) {
               select {
                  do my_ip_c::op1;
                  do my_ip_c::op2;
               }
            }
         }
      }
   };
};
```

Example 332—Use of resource objects as executor claims

Copyright © 2026 Accellera. All rights reserved.

21.7.2.5 Executor query function

The function **executor()** returns a reference to the executor instance currently operative. When called during the evaluation of **exec** blocks of an **action** or flow/resource object or of any function invoked by them, it returns the executor instance assigned to that entity. The function **executor()** can be used, among other purposes, to delegate generic target functions to an executor-specific implementation.

```
function ref executor_base_c executor();
```

Syntax 130—Executor query function

Note that the reference returned from **executor()** for actions assigned to different executors would be different, even if these actions are executing concurrently. The returned value shall be **null** if the evaluating entity is not assigned to any executor. Since assignment to executors is only resolved as part of the solve process, calling **executor()** in **pre_solve exec** blocks shall always return **null**.

In Example 333, a call to the global function `my_target_op()` is delegated to the instance function `my_target_op_impl()` of the currently operative executor, through a call to **executor()**. Function `my_target_op_impl()` is declared in **component executor_base_c** and implemented differently in two executor subtypes. Consequently, the call to `my_target_op()` in the **exec body** of **action** `call_op` will be implemented differently based on the executor assignment of `call_op`.

Copyright © 2026 Accellera. All rights reserved.

```
function void my_target_op(int param) {
   if (executor() != null ) {
      executor().my_target_op_impl(param);
   } else {
      // default implementation
   }
}

extend component executor_base_c {
   function void my_target_op_impl(int param);
};

component A_executor_c : executor_c<> {
   function void my_target_op_impl(int param) {
      // implementation for execution agent of type A
   }
};

component B_executor_c : executor_c<> {
   function void my_target_op_impl(int param) {
      // implementation for execution agent of type B
   }
};

component pss_top {
   executor_group_c<> exe_g;
   A_executor_c a_exe;
   B_executor_c b_exe;

   exec init_down {
      exe_g.add_executor(a_exe);
      exe_g.add_executor(b_exe);
   }

   action call_op {
      rand executor_claim_s<> my_exe;
      exec body {
         my_target_op(10);
      }
   };
};
```

Example 333—Function delegation to executor

21.7.2.6 Executor assignment for components

Each component instance that has defined target exec blocks (*run_start*, *run_end*, *header*, and *declaration*) shall be associated with an executor in which those exec blocks are generated and executed.

```
function void set_executor(ref executor_base_c xtr);
```

Syntax 131—Set_executor function

Copyright © 2026 Accellera. All rights reserved.

A component may explicitly specify its executor by calling the *set_executor* function in its exec init_up or init_down block. Each component instance has at most one associated executor. Consecutive invocation of the *set_executor* function overrides the previous value.

If a component instance does not explicitly assign an executor, it is automatically assigned the executor of its parent component in the component-instance hierarchy. If no executor can be determined for a component instance by explicit assignment or inheritance, and the component type defines one or more component-level target exec blocks, it is an error.

The component's executor serves as the default for actions scheduled on that instance if those actions do not explicitly claim an executor. Actions may claim a different executor explicitly; this does not change the component's executor.

An executor implicitly serves as its own executor and does not require explicit assignment. Assigning any other executor to an executor instance shall be an error.

It shall be an error for an execution unit to explicitly or implicitly assign an executor that references a different execution unit.

21.8 Target execution units

21.8 defines how executors are used by PSS tools to determine how scenario entities are realized. Target execution units are used by PSS tools to group executor realizations into common files that share the same target language, enabling those realizations to be compiled together using the target platform's compilation tools. Each executor instance holds a reference to a target execution unit instance, which tools can use to determine the file in which the executor's realizations should be placed.

```
package executor_pkg {
  enum target_language_e {C, CPP, SV};
  component target_execution_unit_c {
    solve function void set_filename(string filename);
    solve function string get_filename();
    solve function void set_target_language(target_language_e lang);
    solve function target_language_e get_target_language();
  }
  component executor_base_c {
    solve function void set_target_execution_unit(
        ref target_execution_unit_c target_execution_unit);
    solve function ref target_execution_unit_c get_target_execution_unit();
  }
}
```

Syntax 132—Target_execution_unit_c

The PSS target execution unit component type definition is part of the `executor_pkg` package. A target execution unit component is strictly a test-realization artifact. It shall be an error to declare in its scope scenario model elements, namely: **action** types, **pool** instances, and pool binding directives.

Copyright © 2026 Accellera. All rights reserved.

The `target_language_e` enum type is used by the tool to contain the languages supported by the tool and can be extended accordingly.

The target execution unit has the following functions:

— `set_filename/get_filename`: Used to set and get a `filename`. The `filename` is a string used by a PSS tool to construct file names. The `filename` value serves as the root for the name of the generated test file in the file system, with a target-language-specific extension and possibly a test prefix prepended. For example, for C, the tool may generate two files (`<filename>.c` and `<filename>.h`).

  If not specified, the tool will generate a name based on the target execution unit instance name.

— `set_target_language/get_target_language`: Used to set and get the target language. The target language is of type `target_language_e`. The target language is used by a PSS tool to determine the target language to realize the generated test case in. The default target language is C.

— A reference to a target execution unit in an executor is set by the user calling the `set_target_execution_unit` function. The tool will use the reference to find the target execution unit to be used when realizing the scenario entities per the executor instance. It shall be an error to fail to provide a reference to a target execution unit for an executor if target execution units are declared in the system.

Example 334 illustrates how target executor units are configured for a dual cluster SoC. One cluster has an embedded RISC-V core and an embedded DSP core, while the second cluster is coreless and driven by a System Verilog bus transactor for both the RISC-V and DSP buses in the testbench.

For the first cluster, given that the RISC-V and DSP cores use different compilers, and therefore, require separate images compiled from different source files, two target execution units are declared, one for each core. Each target execution unit is assigned a different `filename` value. The code generated for the first cluster places all RISC-V mapped actions in the file associated with the RISC-V target execution unit, and all DSP mapped actions in the file associated with the DSP target execution unit.

For the second cluster, all actions are mapped to SystemVerilog, which only needs a single target execution unit, assuming that all SystemVerilog code in the testbench is compiled by the same tool.

In addition, each cluster has inbound traffic coming from four UART SystemVerilog transactor instances. The actions driving these transactors will be mapped to SystemVerilog code, and therefore, will also be associated with the SystemVerilog target execution unit instantiated as `sv_testbench`.

Copyright © 2026 Accellera. All rights reserved.

```
// Uart executor trait
struct uart_xtor_trait_s : executor_trait_s {
  rand int id;
  rand int cluster_id;
}

// cpu executor trait
struct cpu_executor_trait : executor_trait_s {
  rand int cluster_id;
  rand int core_id;
  rand core_type_e core_type;
};

// executor types
component cpu_executor_c : executor_c<cpu_executor_trait> {}
component c_cpu_executor_c : cpu_executor_c {}
component sv_cpu_executor_c : cpu_executor_c {}
component uart_xtor_executor_c : executor_c<uart_xtor_trait_s> {}


// riscv target execution unit
component riscv_execution_unit_c : target_execution_unit_c {
  exec init_down {
    set_target_language(target_language_e::C);
    set_filename("riscv");
  }
}

// dsp target execution unit
component dsp_execution_unit_c : target_execution_unit_c {
  exec init_down {
    set_target_language(target_language_e::C);
    set_filename("dsp");
  }
}

// sv target execution unit, used for core bus in second cluster
// and all 8 uart transactor stimulus code
component sv_execution_unit_c : target_execution_unit_c {
  exec init_down {
    set_target_language(target_language_e::SV);
    set_filename("sv_testbench");
  }
}
```

Example 334—Dual-cluster SOC example

Copyright © 2026 Accellera. All rights reserved.

```
  component pss_top {
    executor_group_c<uart_xtor_trait_s> uart_xtor_group;
    uart_xtor_executor_c uart_xtor_ex[8];
    c_cpu_executor_c c_cpu_ex[2];
    sv_cpu_executor_c sv_cpu_ex[2];

    executor_group_c<cpu_executor_trait> cpu_group;
    // Target execution units
    riscv_execution_unit_c riscv;
    dsp_execution_unit_c dsp;
    sv_execution_unit_c sv_testbench;

    exec init_down {
      foreach (uart_xtor_ex[i]) {
        uart_xtor_ex[i].set_target_execution_unit(sv_testbench);
        uart_xtor_ex[i].trait.id = i;
        uart_xtor_ex[i].trait.cluster_id = i/4;
        uart_xtor_group.add_executor(uart_xtor_ex[i]);
      }
      c_cpu_ex[0].set_target_execution_unit(riscv);
      c_cpu_ex[0].trait.cluster_id = 0;
      c_cpu_ex[0].trait.core_id = 0;
      c_cpu_ex[0].trait.core_type = RISCV;
      cpu_group.add_executor(c_cpu_ex[0]);

      c_cpu_ex[1].set_target_execution_unit(dsp);
      c_cpu_ex[1].trait.cluster_id = 0;
      c_cpu_ex[1].trait.core_id = 1;
      c_cpu_ex[1].trait.core_type = DSP;
      cpu_group.add_executor(c_cpu_ex[1]);

      sv_cpu_ex[0].set_target_execution_unit(sv_testbench);
      sv_cpu_ex[0].trait.cluster_id = 1;
      sv_cpu_ex[0].trait.core_id = 0;
      sv_cpu_ex[0].trait.core_type = RISCV;
      cpu_group.add_executor(sv_cpu_ex[0]);

      sv_cpu_ex[1].set_target_execution_unit(sv_testbench);
      sv_cpu_ex[1].trait.cluster_id = 1;
      sv_cpu_ex[1].trait.core_id = 0;
      sv_cpu_ex[1].trait.core_type = DSP;
      cpu_group.add_executor(sv_cpu_ex[1]);
    }
  }
```

Example 334—Dual-cluster SOC example (cont.)

21.9 Synchronization and communication

Concurrently-executing PSS test realization code may need to dynamically synchronize or communicate at runtime. The PSS core library provides constructs for capturing target-time synchronization and communication such that PSS processing tools can implement that behavior in a portable manner.

PSS defines synchronization and communication constructs in the *sync_pkg* package. In subsequent sections, the enclosing *sync_pkg* is omitted for brevity.

Copyright © 2026 Accellera. All rights reserved.

21.9.1 Channel type

A PSS *channel* is a component that supports target-time communication. All functions within the component are *target* functions. The element type of a channel is specified via its *T* parameter. A channel can hold up to `DEPTH` items (default 1). Data is stored in FIFO order, such that the first call to `get` returns the first data item stored with `put`.

```
component channel_c<type T, int DEPTH=1> {
  target function T get();
  target function void put(T t);
  target function bool try_get(output T t);
  target function bool try_put(T t);
}
```

Syntax 133—PSS channel

The following also apply:

a) Channels can only hold the following data types: numeric types, Booleans, enumerated types that have a base type, packed structs, and arrays thereof.

b) If `DEPTH` is specified, it shall be positive.

c) Data is copied in/out of the channel by value.

d) Implementations shall implement mutual-exclusion and synchronization mechanisms to support data passing between test realization running on multiple executors independent of the language implementing those executors.

21.9.1.1 Get

The *get* method blocks the caller until an element is available in the channel and returns the oldest element. The *get* method is thread-safe, such that if multiple threads call *get* simultaneously an element within the channel is only returned to a single thread. The order in which data is returned to multiple waiters is non-deterministic. See 20.8 for more details on test realization cooperative multitasking that applies when a *get* call blocks.

21.9.1.2 Put

The *put* method blocks the caller until space is available within the channel and inserts the data element into the channel. The *put* method is thread-safe, such that if multiple threads call *put* simultaneously, no data will be lost. When multiple threads call *put* simultaneously, the order in which data is inserted is non-deterministic. See 20.8 for more details on test realization cooperative multitasking that applies when a *put* call blocks.

21.9.1.3 Try_get

The *try_get* method attempts to acquire the oldest element in the channel. If the channel contains at least one element, the oldest element is returned via the output parameter and the function returns *true*. If the channel contains no elements, the function returns *false*.

The *try_get* method is thread-safe, such that a single channel element is only returned to a single simultaneous caller. If the number of elements in the channel is greater than or equal to the number of simultaneous callers (N), then *try_get* will return *true* to all N callers.

Copyright © 2026 Accellera. All rights reserved.

21.9.1.4 Try_put

The *try_put* method attempts to insert a data element into the channel. If the channel has space for a new element, the *try_put* function inserts the element and returns *true*. If the channel does not have space, the *try_put* function returns *false*.

The *try_put* method is thread-safe, such that no data is lost or corrupted in the presence of multiple simultaneous callers.

21.10 Address spaces

The *address space* concept is introduced to model memory and other types of storage in a system. An address space is a space of storage atoms accessible using unique addresses. System memory, external storage, internal SRAM, routing tables, memory mapped I/O, etc., are entities that can be modeled with address spaces in PSS.

An address space is composed of *regions*. Regions are characterized by user-defined properties called *traits*. For example, a trait could be the type of system memory of an SoC, which could be DRAM or SRAM. *Address claims* can be made by scenario entities (actions/objects) on an address space with optional constraints on user-defined properties. An *address space handle* is an opaque representation of an address within an address space.

Standard operations are provided to read data from and write data to a byte-addressable address space. *Registers* and *register groups* are allocated within an address space and use address space regions and handles to read and write register values. Data layout for packed PSS structs is defined for byte-addressable address spaces.

The PSS built-in package **`addr_reg_pkg`** defines types and functions for registers, address spaces, address allocation and operations on address spaces. In subsequent sections, except Syntax 134, the enclosing **`addr_reg_pkg`** is omitted for brevity. Examples may also omit import of **`addr_reg_pkg`** and **`std_pkg`**.

21.10.1 Address space categories

This section describes address space categories.

21.10.1.1 Base address space type

An *address space* is a set of storage atoms accessible using unique addresses. Actions/objects may allocate one or more atoms for their exclusive use.

Address spaces are declared as **components**. **`addr_space_base_c`** is the base type for all other address space types. This component cannot be instantiated directly. The definition of **`addr_space_base_c`** is shown in Syntax 134.

```
package addr_reg_pkg {
  component addr_space_base_c {};
...
}
```

Syntax 134—Generic address space component

Copyright © 2026 Accellera. All rights reserved.

21.10.1.2 Contiguous address spaces

A *contiguous address space* is an address space whose addresses are non-negative integer values. and whose atoms are contiguously addressed. Multiple atoms can be allocated in one contiguous chunk.

Byte-addressable system memory and blocks of data on disk drive are examples of contiguous address spaces.

A contiguous address space is defined by the built-in library component **`contiguous_addr_space_c`** shown in Syntax 135 below. The meanings of the struct type **`addr_trait_s`** and the template parameter **`TRAIT`** are defined in 21.10.2. Address space regions are described in 21.10.3.

```
struct addr_trait_s {};

struct empty_addr_trait_s : addr_trait_s {};

typedef chandle addr_handle_t;

component contiguous_addr_space_c <struct TRAIT : addr_trait_s =
  empty_addr_trait_s> : addr_space_base_c
{
  solve function addr_handle_t add_region(addr_region_s <TRAIT> r);
  solve function addr_handle_t add_nonallocatable_region(addr_region_s <> r);

  bool byte_addressable = true;
};
```

Syntax 135—Contiguous address space component

A contiguous address space is created in a PSS model by creating an instance of **component** **`contiguous_addr_space_c`** in a top-level **component** or any other **component** instantiated under the top-level **component**.

21.10.1.2.1 Add_region function

The **`add_region`** function of contiguous address space components is used to add allocatable address space regions to a contiguous address space. The function returns an address handle corresponding to the start of the region in the address space. Actions and objects can allocate space only from allocatable regions of an address space.

Address space regions are defined in 21.10.3. Address space regions are part of the static component hierarchy. The **`add_region`** function may only be called in **exec init_down** and **init_up** blocks. Address handles are defined in 21.13.3.

21.10.1.2.2 Add_nonallocatable_region function

The **`add_nonallocatable_region`** function of contiguous address space components is used to add non-allocatable address space regions to a contiguous address space. The function returns an address handle corresponding to the start of the region in the address space.

Copyright © 2026 Accellera. All rights reserved.

The address space allocation algorithm shall not use non-allocatable regions for allocation.

Address space regions are defined in 21.10.3. Address space regions are part of the static component hierarchy. The **`add_nonallocatable_region`** function may only be called in **exec init_down** and **init_up** blocks. Address handles are defined in 21.13.3.

21.10.1.2.3 Example

Example 335 demonstrates instantiating an address space and adding regions to it (for the definition of **struct `addr_region_s`**, see 21.10.3.2).

```
component pss_top {

   import addr_reg_pkg::*;

   my_ip_c ip;

   contiguous_addr_space_c<> sys_mem;

   exec init_up {
          // Add regions to space here
          addr_region_s<> r1;
          r1.size = 0x40000000; // 1 GB
          (void)sys_mem.add_region(r1);

          addr_region_s<> mmio;
          mmio.size = 4096;
          (void)sys_mem.add_nonallocatable_region(mmio);
   }
}
```

Example 335—Contiguous address space in pss_top

21.10.1.3 Byte-addressable address spaces

A *byte-addressable* space is a contiguous address space whose storage atom is a byte and to/from which PSS data can be written/read using standard generic operations. The PSS core library standardizes generic APIs to write data to or read data from any address value as bytes. The read/write API and data layout of PSS data into a byte-addressable space are defined in 21.13.

By default, **component `contiguous_addr_space_c`** is a byte-addressable space unless the **`byte_addressable`** Boolean field is set to *false*.

21.10.1.4 Transparent address spaces

Transparent address spaces are used to enable transparent claims—constraining and otherwise operating on concrete address values on the solve platform. For more information on transparent address claims, see 21.11.4.

All regions of a transparent space provide a concrete start address and the size of the region. Only transparent regions (see 21.10.3.3) may be added to a transparent address space using function **`add_region()`**. Note however that transparent regions may be added to a non-transparent space.

Component `transparent_addr_space_c` is used to create a transparent address space (see Syntax 136). See Example 337.

Copyright © 2026 Accellera. All rights reserved.

```
component transparent_addr_space_c
                    <struct TRAIT: addr_trait_s = empty_addr_trait_s>
                    : contiguous_addr_space_c<TRAIT> {};
```

Syntax 136—Transparent address space component

21.10.1.5 Other address spaces

Other kinds of address spaces, with different assumptions on allocations and generic operations, are possible. These may be represented as derived types of the corresponding base space/region/claim types. An example could be a space representing a routing table in a network router. PSS does not attempt to standardize these.

21.10.2 Address space traits

An address space *trait* is a PSS **struct**. A trait **struct** describes properties of a contiguous address space and its regions. **`empty_addr_trait_s`** is defined as an empty trait struct that is used as the default trait type for address spaces, regions and claims.

All regions of an address space share a trait *type*. Every region has its specific trait *value*.

```
package ip_pkg {

  struct mem_trait_s : addr_trait_s {
    rand mem_kind_e        kind;
    rand cache_attr_e      ctype;
    rand int in [0..3]     sec_level;
    rand bool              mmio;
  };

};
```

Example 336—Example address trait type

Copyright © 2026 Accellera. All rights reserved.

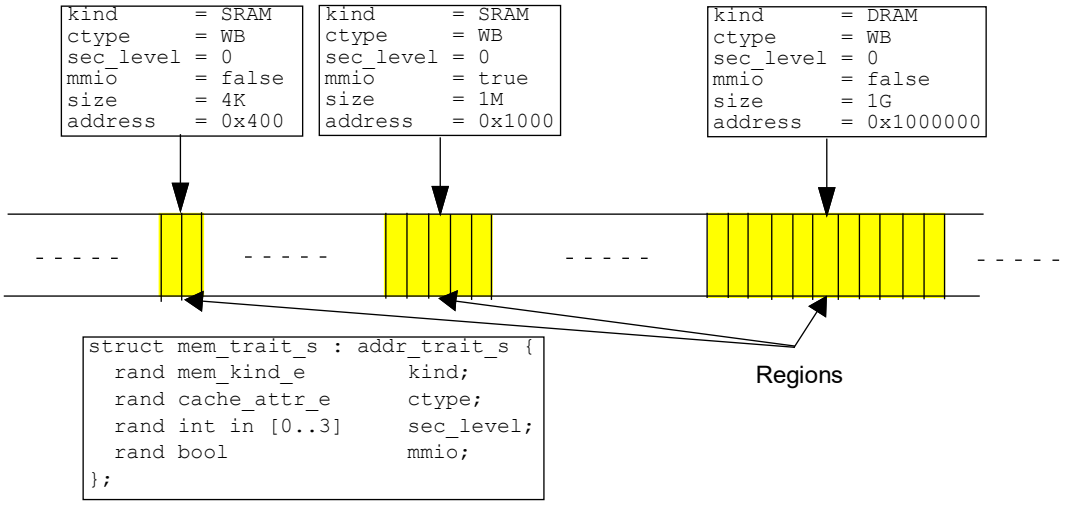


Figure 52—Address space regions with trait values

```
component pss_top {

   import addr_reg_pkg::*;
   import ip_pkg::*;

   // IP component
   my_ip_c ip;

   // mem_trait_s trait struct is used for sys_mem address space
   transparent_addr_space_c<mem_trait_s> sys_mem;

   exec init_up {
      // Add regions to space here. All regions added to sys_mem space
      // shall have trait type mem_trait_s

      transparent_addr_region_s<mem_trait_s> sram_region;

      sram_region.trait.kind      = SRAM;
      sram_region.trait.ctype     = WB;
      sram_region.trait.sec_level = 0;
      sram_region.trait.mmio      = false;
      sram_region.size            = 4096;
      sram_region.addr            = 0x400;

      (void)sys_mem.add_region(sram_region);

      // add other regions
      // ...
    }
}
```

Example 337—Address space with trait

Copyright © 2026 Accellera. All rights reserved.

21.10.3 Address space regions

An address space may be composed of *regions*. Regions map to parts of an address space. A region may be characterized by values assigned to address space traits. Traits define properties of a region. Specific constraints are placed on *address claim* traits to allocate addresses from regions with desired characteristics. Regions with trait values that satisfy the claim's trait constraints are the candidate matching regions. An address claim may span more than one region that satisfies claim trait constraints.

Address space regions are part of the static component hierarchy. The **`add_region`** and **`add_nonallocatable_region`** functions (see 21.10.1.2.1 and 21.10.1.2.2) may only be called in **exec init_down** and **init_up** blocks.

21.10.3.1 Base region type

`addr_region_base_s` is the base type for all address space regions (see Syntax 137). Specifying a value for the **`size`** field is required. Specifying a value for the **`tag`** field is optional.

```
struct addr_region_base_s {
  bit[64] size;
  string  tag;
};
```

Syntax 137—Base address region type

The **`tag`** associated with the region from which a memory claim is satisfied may be retrieved using the **`get_tag()`** function (see 21.13.8).

21.10.3.2 Contiguous address regions

The *addr_region_s* type represents a region in contiguous address space (see Syntax 138). The region type is fully characterized by the template **`TRAIT`** parameter value and the **`size`** attribute of the base region type.

```
struct addr_region_s <struct TRAIT : addr_trait_s = empty_addr_trait_s>
                     : addr_region_base_s {
  TRAIT trait;
};
```

Syntax 138—Contiguous address space region type

The values of the trait struct attributes describes the contiguous address region. The PSS tool will match the trait attributes of regions to satisfy an address claim as described in 21.11. See an example of trait attribute setting in 21.11.8.

21.10.3.3 Transparent address regions

The **`transparent_addr_region_s`** type defines a *transparent* region over a contiguous address space. *Transparent* means that the region's start (lower) address is known to the PSS tool for solve-time resolution of a claim address within the address space.

The **`addr`** field of this region is assigned the start address of the region. The end address of the region is the calculated value of the expression: **`addr`** + **`size`** - 1.

Copyright © 2026 Accellera. All rights reserved.

See Example 337 where a transparent region is added to a transparent address space.

```
struct transparent_addr_region_s
        <struct TRAIT : addr_trait_s = empty_addr_trait_s>
        : addr_region_s<TRAIT> {
  bit[64] addr;
};
```

Syntax 139—Transparent region type

21.11 Allocation within address spaces

The PSS input model can allocate storage atoms from an address space. Allocation can be made for an exclusive use by a behavior or for a shared use by multiple behaviors. An example of an exclusive use is a DMA controller action allocating a buffer in system memory for output data. An example of a shared use is a cache false sharing compound action allocating a cache line to be accessed by concurrent write/read cache **action**s.

All address space allocations are done in the declarative domain of a PSS input model. An *address claim struct*, defined in the following sections, is used for allocation.

An instance of an address claim struct describes an address claim on an address space. A claim is *matched* to the address space nearest in the **component** instance tree, whose trait type matches the claim trait type (see 21.11.7). A claim is satisfied by allocation from a region (or regions) whose trait value satisfies the constraints on the claim trait (see 21.11.5).

A claim struct can be instantiated under an **action**, a flow object or resource object, or any of their nested structs. The declaration of a claim struct instance causes allocation to occur when the declaring object is instantiated or the **action** is traversed.

21.11.1 Base claim type

The **`addr_claim_base_s`** struct (see Syntax 140) is the base type for all address space claims.

```
struct addr_claim_base_s {
  rand bit[64] size;
  rand bool permanent;
  constraint default permanent == false;
};
```

Syntax 140—Base address space claim type

21.11.2 Contiguous claims

An address claim can be made on a contiguous address space by declaring a **struct** of type **`addr_claim_s`**. This claim is also known as an *opaque* claim. The absolute address of the claim is not assumed to be known at solve time.

This standard does not define any method by which the PSS tool might resolve address claims at solve time or might generate code for runtime allocation. One possible method could be PSS tool-specific APIs for

Copyright © 2026 Accellera. All rights reserved.

solve-time and runtime allocation. The *address space handle* obtained from a claim shall fall within a region or regions whose traits satisfy the claim constraints.

An address claim in contiguous address space is always a contiguous chunk of addresses, potentially spanning multiple regions that are adjacent.

An address claim can be made on transparent (described below, in 21.11.4) or non-transparent address spaces.

```
struct addr_claim_s
                < struct TRAIT : addr_trait_s = empty_addr_trait_s,
                  struct ALLOC_MODE : alloc_base_mode_s = alloc_base_mode_s>
                : addr_claim_base_s {
  rand TRAIT trait;
  rand bit[64] in [64'd2**0, 64'd2**1, 64'd2**2, 64'd2**3 , 64'd2**4 ,
    64'd2**5 , 64'd2**6 , 64'd2**7 , 64'd2**8 , 64'd2**9 , 64'd2**10,
    64'd2**11, 64'd2**12, 64'd2**13, 64'd2**14, 64'd2**15, 64'd2**16
    64'd2**17, 64'd2**18, 64'd2**19, 64'd2**20, 64'd2**21, 64'd2**22
    64'd2**23, 64'd2**24, 64'd2**25, 64'd2**26, 64'd2**27, 64'd2**28
    64'd2**29, 64'd2**30, 64'd2**31, 64'd2**32, 64'd2**33, 64'd2**34
    64'd2**35, 64'd2**36, 64'd2**37, 64'd2**38, 64'd2**39, 64'd2**40
    64'd2**41, 64'd2**42, 64'd2**43, 64'd2**44, 64'd2**45, 64'd2**46
    64'd2**47, 64'd2**48, 64'd2**49, 64'd2**50, 64'd2**51, 64'd2**52
    64'd2**53, 64'd2**54, 64'd2**55, 64'd2**56, 64'd2**57, 64'd2**58
    64'd2**59, 64'd2**60, 64'd2**61, 64'd2**62, 64'd2**63] alignment;
  rand ALLOC_MODE alloc_mode;
};
```

Syntax 141—Contiguous address space claim type

The **`alignment`** attribute specifies the address alignment of the resolved claim address.

The **`alloc_mode`** attribute specifies the allocation mode (see 21.11.3) of the resolved claim address. **`alloc_mode`** attributes are used as inputs for the address resolution phase and **`alloc_mode`** attributes are resolved during the solve phase.

21.11.3 Allocation mode

The allocation mode struct captures address sharing relationships with other claims as well as optional directives to memory allocation mechanisms. By default, all claims are exclusive, and addresses will be allocated for each claim in a scenario.

Attributes of the allocation mode struct can be used to override this default: different claims can share the same start address, share overlapping addresses, some claims can skip allocation, etc.

Copyright © 2026 Accellera. All rights reserved.

```
enum alloc_access_mode_e {EXCLUSIVE, SHARED};
enum alloc_share_mode_e {RANDOM, OVERLAP, SAME_ADDRESS};
enum alloc_skip_mode_e {DONT_SKIP, SKIP};

struct alloc_base_mode_s {
  rand alloc_access_mode_e access;
  constraint default access == EXCLUSIVE;
  rand alloc_share_mode_e sharing;
  constraint default sharing == RANDOM;
  rand alloc_skip_mode_e skipping;
  constraint default skipping == DONT_SKIP;
  rand int tag;
  constraint default tag == -1;
};
```

Syntax 142—Allocation mode type

The access attribute specifies whether the claim is sharing with other claims or whether the claim is exclusive. The access attribute is of enum type **`alloc_access_mode_e`**.

Type **`alloc_access_mode_e`** has the following enum items:

- EXCLUSIVE—**`alloc_mode`**'s default value. Specifies that a claim is exclusive within its available storage atoms.
- SHARED—shared claim. The claim can overlap with another shared claim in its available storage atoms.

The sharing attribute specifies how to share storage atoms with other claims. The sharing attribute is of type **`alloc_share_mode_e`**.

Type **`alloc_share_mode_e`** has the following enum items:

- OVERLAP—shared claim. The claim's storage atoms shall overlap with storage atoms of all shared claims that have shared mode OVERLAP with the same tag value. The tag value shall be greater or equal to zero. If there is no shared claim that has shared mode OVERLAP with the same tag, the claim is treated as RANDOM. A shared claim that has shared mode OVERLAP can share storage atoms with other SHARED claims (e.g., SHARED/SAME_ADDRESS claim). A shared claim that has shared mode OVERLAP with a matching tag to a claim with a shared mode of SAME_ADDRESS shall overlap with that claim.
- SAME_ADDRESS—shared claim. The claim shall have the same address as all shared claims that are SAME with the same tag value within its available address space. The tag value shall be greater or equal to zero. Enables relating addresses based on some criteria. If there is no shared claim with the same tag, the claim is treated as RANDOM.
- RANDOM—shared claim, but there is no specific additional requirement. The claim can share in any way the storage atoms of another SHARED claim, or it can even be exclusive.

The skipping attribute specifies whether the claim is to be allocated (not skipped) or whether the claim is not to be allocated (skipped). It enables the user to allocate a claim conditionally.

Type **`alloc_skip_mode_e`** has the following enum items:

Copyright © 2026 Accellera. All rights reserved.

— DONT_SKIP—the default skipping value is not to skip. The memory allocation mechanism will process this claim.
— SKIP—voids the claim and therefore the memory allocation mechanism can ignore the claim. This feature is needed to enable conditional allocation of address claims

Example 338 illustrates how a shared claim is made for memory to store an embedded core firmware image. Core executors in the same compilation unit share the same image. An address claim is done within the executor's resource. Executors within the same cluster should claim the same address and size for the cluster's firmware image that will be loaded in memory. The tag is used to enable the address constraint solver to pick the same address for any two executors in the same cluster. The tag is constrained to be equivalent to the cluster's unique ID (see Clause 9). The memory claim trait needs to be the same for all executors in the same cluster. The claim is voided if the executor target is not an embedded core (i.e., `trait.lang != C`) by constraining the skipping attribute to SKIP.

```
package user_executor_pkg {
  import executor_pkg::*;
  import addr_reg_pkg::*;
  import addr_pkg::*;

  const int NOF_EXECUTORS = 4;
  const int NOF_CLUSTERS = 2;
  enum lang_e {C, SV};

  struct cpu_executor_trait : executor_trait_s {
    rand int in [0..NOF_EXECUTORS-1] cpu_id;
    rand int in [0..NOF_CLUSTERS-1] cluster_id;
    rand lang_e lang;
  };

  resource executor_r : executor_claim_s < cpu_executor_trait > {
    constraint instance_id == trait.cpu_id;
    rand addr_claim_s<mem_addr_trait_s> claim;
    constraint claim.size == 128;
    constraint default disable claim.alloc_mode.access;
    constraint default disable claim.alloc_mode.sharing;
    constraint default disable claim.alloc_mode.skipping;
    constraint
      if (trait.lang == C) {
        claim.alloc_mode.access == SHARED;
        claim.alloc_mode.skipping == DONT_SKIP;
        claim.alloc_mode.sharing == SAME_ADDRESS;
      } else {
       claim.alloc_mode.skipping == SKIP;
      };
    constraint default disable claim.alloc_mode.tag;

    // assuming pss_top.cluster is an array of PSS component instances
    // representing the available clusters;
    constraint
       claim.alloc_mode.tag == pss_top.cluster[trait.cluster_id].uid;
  };
}
```

Example 338—Shared claim with same address for embedded code using same compilation unit

Copyright © 2026 Accellera. All rights reserved.

Example 339 illustrates how overlapping shared claims are specified. The address space has four regions with 1024 bytes each. The `big_claim_a` action randomly picks a region. The `big_claim_a` constrains allocation for four `share_addr_a` actions: 1) Constrain for allocation of size 8 bytes each (constraining size attribute). 2) Constrain for sharing bytes (using SHARED allocation access mode). 3) Constrain to have the same traits. 4) Constrain to how sharing between the actions shall overlap. 1 and 2 shall overlap and 2 and 3 shall overlap, based on the constraint over their respective tag attribute. The chosen tag is the unique ID of one of the two action handles.

```
component pss_top {
  action share_addr_a {
    rand transparent_addr_claim_s<trait_s>  addr_claim;
    constraint default disable addr_claim.alloc_mode.access;
    constraint addr_claim.alloc_mode.access == SHARED;
    constraint default disable addr_claim.alloc_mode.sharing;
    constraint addr_claim.alloc_mode.sharing == OVERLAP;
    exec post_solve {
      print("TAG : %0d\n", addr_claim.alloc_mode.tag);
    }
  }
  action big_claim_a {
    share_addr_a s[4];
    rand trait_s trait;
    constraint trait.region_id in [0..3];
    // Claim from action s[0] shall overlap with claim from action s[1]
    // and vice-versa
    constraint s[0].addr_claim.alloc_mode.tag == s[1].uid;
    constraint s[1].addr_claim.alloc_mode.tag == s[1].uid;
    // Claim from action s[2] shall overlap with claim from action s[3]
    // and vice-versa
    constraint s[2].addr_claim.alloc_mode.tag ==
      s[3].uid;
    constraint s[3].addr_claim.alloc_mode.tag ==
      s[3].uid;

    activity {
      parallel {
        replicate (i:4)
          s[i] with {
            addr_claim.size == 8;
            addr_claim.trait == this.trait;
            default disable addr_claim.alloc_mode.tag;
          }
        }
      }
    }
  }
}
```

Example 339—Overlapping shared claims

21.11.4 Transparent claims

A claim of type **`transparent_addr_claim_s`** (see Syntax 143) is required to make a transparent claim on a transparent contiguous address space. A transparent claim is characterized by the absolute allocation address attribute (**`addr`**) of the claim. A transparent claim is associated with the nearest address space with the same trait type, in the same way that a non-transparent claim is. However, a transparent claim

Copyright © 2026 Accellera. All rights reserved.

that is thereby associated with a non-transparent space shall be flagged as an error. The PSS tool has all the information at solve time about the transparent address space necessary to perform allocation within the limits of the address space. More details about allocation and claim lifetime can be found in the following section.

The **addr** field of this claim type can be used to put a constraint on an absolute address of a claim.

```
struct transparent_addr_claim_s
  <struct TRAIT : addr_trait_s = empty_addr_trait_s,
   struct ALLOC_MODE : alloc_base_mode_s = alloc_base_mode_s > :
   addr_claim_s<TRAIT, ALLOC_MODE> {
  rand bit[64] addr;
}
```

Syntax 143—Transparent contiguous address space claim type

Example 340 illustrates how a transparent claim is used. A transparent address claim is used in **action** `my_op`. A constraint is placed on the absolute resolved address of the claim. This is possible only because of the transparent address space that contain transparent regions where the base address of the region is known at solve time.

```
component pss_top {

  transparent_addr_space_c<> mem;

  action my_op {
    rand transparent_addr_claim_s<> claim;
    constraint claim.size == 20;

    // Constraint on absolute address
    constraint (claim.addr & 0x3) == 0x1;
  };

  exec init_up {
    transparent_addr_region_s<> region1, region2;
    region1.size = 50;
    region1.addr = 0x10000;
    (void)mem.add_region(region1);

    region2.size = 10;
    region2.addr = 0x20000;
    (void)mem.add_region(region2);
  }
};
```

Example 340—Transparent address claim

21.11.5 Claim trait semantics

Constraints placed on the trait attribute of a claim instance shall be satisfied by the allocated addresses. Allocated addresses shall be in regions whose trait values satisfy claim trait constraints.

See an example in 21.11.8.

Copyright © 2026 Accellera. All rights reserved.

21.11.6 Allocation consistency

An address claim struct is resolved to represent the allocation of a set of storage atoms from the nearest storage space, for the exclusive use of actions that can access the claim attribute. In the case of a contiguous address space, the set is a contiguous segment, from the start address to the start address + **`size`** - 1. All addresses in the set are uniquely assigned to that specific instance of the address claim struct for the duration of its lifetime, as determined by the actions that can access it (see details below). Two instances of an address claim struct shall resolve to mutually exclusive sets of addresses if

— Both are taken from the same address space, and

— An action that has access to one may overlap in execution time with an action that has access to the other.

The number of storage atoms in an allocation is represented by the attribute **`size`**.

The start address is represented directly by the attribute **`addr`** in **`transparent_addr_claim_s<>`**, or otherwise obtained by calling the function **`addr_value()`** on the address space handle returned by **`make_handle_from_claim()`**.

Following is the definition of the lifetime of scenario entities:

Table 32—Scenario entity lifetimes

Entity	Lifetime
Atomic action	From the time of **exec body** entry (immediately before executing the first statement) to the time of the **exec body** exit (immediately after executing the last statement).
Compound action	From the start time of the first sub-action(s) to the end time of the last sub-action(s).
Flow object	From the start time of the action outputting it (for the initial state, the start time of the first action in the scenario) to the end time of the last action(s) inputting it (if any) or the end-time of the last action outputting it (if no action inputs it).
Resource object	From the start time of the first action(s) locking/sharing it to the end time of the last action(s) locking/sharing it.
Struct	Identical with the entity that instantiates it.

The lifetime of the allocation to which a claim struct resolves, and hence the exclusive use of the set of addresses, may be extended beyond the scenario entity in which the claim is instantiated in one of two ways:

— A handle that originates in a claim is assigned to entities that have no direct access to the claim in solve execs (for definition of address space handles, see 21.13.3). For example, if an action assigns a handle field (of type **`addr_handle_t`**) of its output **buffer** object with a handle it obtained from its own claim, the allocation lifetime is extended to the end of the last action that inputs that **buffer** object.

— The attribute **`permanent`** is constrained to *true*, in which case the lifetime of the claim is extended to the end of the test.

21.11.6.1 Example

The example below demonstrates how the scheduling of actions affects possible resolutions of address claims. In this model, **action** `my_op` claims 20 bytes from an address space, in which there is one region of size 50 bytes and another of size 10. In **action** `test1,` the three **action**s of type `my_op` are scheduled

Copyright © 2026 Accellera. All rights reserved.

sequentially, as the iterations of a **repeat** statement. No execution of `my_op` overlaps in time with another, and therefore each one can be allocated any set of consecutive 20 bytes, irrespective of previous allocations. Note that all three allocations shall come from the 50-byte region, as the 10-byte region cannot fit any of them. In `test2`, by contrast, the three **action**s of type `my_op` expanded from the **replicate** statement are scheduled in parallel. This means that they would overlap in execution time, and therefore need to be assigned mutually exclusive sets of addresses. However, such allocation is not possible out of the 50 bytes available in the bigger region. Here too, the smaller region cannot fit any of the three allocations. Nor can it fit part of an allocation, because it is not known to be strictly contiguous with the other region.

```
component pss_top {
  action my_op {
    rand addr_claim_s<> claim;
    constraint claim.size == 20;
  };

  contiguous_addr_space_c<> mem;

  exec init_up {
    addr_region_s<> region1, region2;
    region1.size = 50;
    (void)mem.add_region(region1);
    region2.size = 10;
    (void)mem.add_region(region2);
  }

  action test1 {
    activity {
      repeat (3) {
        do my_op; // OK – allocations can be recycled
      }
    }
  };

  action test2 {
    activity {
      parallel {
        replicate (3) {
          do my_op; // error – cannot satisfy concurrent claims
        }
      }
    }
  };
};
```

Example 341—Address space allocation example

Copyright © 2026 Accellera. All rights reserved.

21.11.7 Rules for matching a claim to an address space

a) A claim is associated with a unique address space based on the static structure of the model.
b) A claim is resolved to an address space that:
    1) matches the trait type of the claim
    2) is instantiated in a containing **component** of the current scenario entity (the context **component** hierarchy of an action or the container **component** of a flow/resource object pool)
    3) is nearest in the **component** hierarchy going up from the context **component** to the root **component**
c) It shall be an error if more than one address space matches a claim at the component context identified in b).

21.11.8 Allocation example

In following example, `pss_top` has instances of the `sub_ip` and `great_ip` components. `sub_ip` is composed of the `good_ip` and `great_ip` components. `good_ip` and `great_ip` allocate space with trait `mem_trait_s`. Memory allocation in the `top_gr_ip` instance of `pss_top` will be matched to the `sys_mem` address space that is instantiated in `pss_top`. Memory claims in `gr_ip` and `go_ip` from `pss_top.sub_system` will be matched to the address space in `sub_ip`, as the `sub_ip` address_space will be the nearest space with a matching trait in the component tree.

Note how within the two address spaces, there are regions with the same base address. Claims from actions of the two instances of `great_ip` may be satisfied with overlapping addresses even if they are concurrent, since they are taken out of different address spaces.

Copyright © 2026 Accellera. All rights reserved.

```
import addr_reg_pkg::*;
import mem_pkg::*;

package mem_pkg {
   enum cache_attr_e {UC, WB, WT, WC, WP};

   struct mem_trait_s : addr_trait_s {
      rand cache_attr_e  ctype;
      rand int in [0..3] sec_level;
   }
};

component good_ip {
   action write_mem {
      // Allocate from nearest address space matching TRAIT type and value
      rand transparent_addr_claim_s<mem_trait_s> mem_claim;

      constraint mem_claim.size == 128;
      constraint mem_claim.trait.ctype == UC;
   }

   action write_mem_unconstrained {
      // Allocate from nearest address space matching TRAIT type and value

      // Note that ctype field of the claim trait is unconstrained.
      // However, given there is only a single region in the address space
      // with ctype==UC, that region is chosen as it is the only match
      // available that can satisfy the trait constraints.

      // ctype cannot be randomized to have a value that is not UC because
      // it is compelled to match with one of the regions, just like when
      // an action wants to consume a buffer object, it needs to pick from
      // the available objects in the pool.
      rand transparent_addr_claim_s mem_claim;
      constraint mem_claim.size == 128;
   }
};

component great_ip {
   action write_mem {

      // Allocate from nearest address space matching TRAIT type and value
      rand transparent_addr_claim_s<mem_trait_s> mem_claim;

      constraint mem_claim.size == 256;
      constraint mem_claim.trait.ctype == UC;
   }
};

component sub_ip {
   // Subsystem has its own address space
   transparent_addr_space_c<mem_trait_s> mem;

   good_ip  go_ip;
   great_ip gr_ip;
};
```

Example 342—Address space allocation example

Copyright © 2026 Accellera. All rights reserved.

```
component pss_top {
   sub_ip sub_system;
   great_ip top_gr_ip;

   transparent_addr_space_c<mem_trait_s> sys_mem;

   exec init_up {
      transparent_addr_region_s<mem_trait_s> region;

      region.size            = 1024;
      region.addr            = 0x8000;
      region.trait.ctype     = UC;
      region.trait.sec_level = 0;

      transparent_addr_region_s<mem_trait_s> great_region;

      great_region.size            = 1024;
      great_region.addr            = 0x8000;
      great_region.trait.ctype     = UC;
      great_region.trait.sec_level = 2;

      (void)sys_mem.add_region(region);

      (void)sub_system.mem.add_region(great_region);
   };
};
```

Example 342—Address space allocation example (cont.)

21.11.9 Address claim resolution

As mentioned in 11.3.1.1, address claim resolution is another “action traversal” step. Address resolution occurs before execution of a **pre_body** *exec block* of their respective actions. The address resolution process is required to meet all the requirements specified in 21.11 for each of the claims.

21.12 Address space group

Different IP PSS models may have different usage models for claiming address space storage atoms. An *address space group* defines the union of multiple individual address spaces that share common storage elements. The PSS input model can allocate common storage elements for the exclusive use of certain behaviors.

The usage model is determined by the address space trait type, the address space region types, etc. Address space group enables the integration of IP PSS models such that each IP PSS model has a different view to common storage atoms.

The component type `addr_space_group_c` is used to group one or more address spaces.

Copyright © 2026 Accellera. All rights reserved.

```
package addr_reg_pkg {
  component addr_space_group_c  {
    function void add_addr_space(ref addr_space_base_c address_space);
  }
}
```

Syntax 144—Address space group

21.12.1 Function add_addr_space

Instance function `add_addr_space` (see Syntax 144) of `addr_space_group_c` is used to populate the group with address space instances. The `add_addr_space` function may only be called in `exec init_down` and `init_up` blocks.

The following also apply:

a) Any address space can be added to a given group, regardless of where it is instantiated in the component instance tree. This includes address spaces instantiated above the group, below it, or in a different subtree.

b) An address space instance may not be added more than once to the same group.

c) An address space instance may not be added to more than one group.

d) An address space does not have to be added to any group. An address space not added to any group will not share storage atoms with other groups and will follow address space semantics mentioned above (will follow standalone address space semantics).

Example 343 demonstrates how two address claim usage models for two different IP PSS models are integrated using an address space group. IP **a** address claim use-model is to get 8 or 16 bytes from an address space with 256 bytes. IP **a** address space has 2 regions. Users can control which region is selected via the trait attribute **id**. IP **b** address claim use-model gets 2 bytes from an address space with 256 bytes. IP **b** address space has 128 regions. Users can control which region is selected via the trait attribute **id**. In `pss_top` there are two address space instances `mema` and `memb`, each using a different trait type. Actions in IP **a** memory claims will match with address space `mema` and actions in IP **b** memory claims will match with address space `memb`. `mema` and `memb` are added to the address space group instance `mem_group`. Both actions share 256 storage atoms. The test case in Example 343 schedules three actions in parallel; therefore, they should all get exclusive storage atoms from the common 256 storage atoms.

Copyright © 2026 Accellera. All rights reserved.

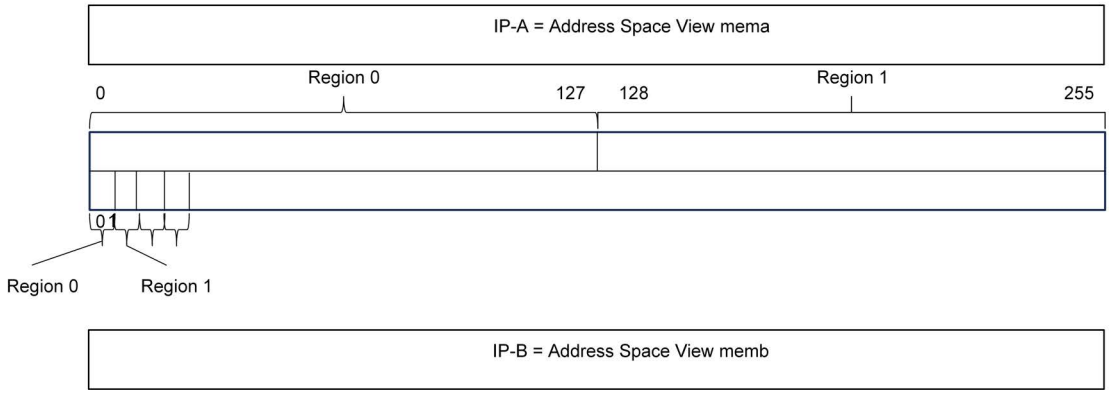


Figure 53—Different IP views of common storage atoms

Figure 53 demonstrates how address claims can be satisfied coming from two IPs using common storage atoms. When claims from two IPs overlap in time, exclusive storage atoms are provided.

Table 33—Overlapping and sequential address claims examples

Time window	A claim bytes	B claim bytes	A address	B address
0	8	2	0x0-0x7	0x8-0x9
1	8		0x8-0xf	
2		2		0x8-0x9
3	16	2	0x4-0x13	0x0-0x1

Copyright © 2026 Accellera. All rights reserved.

```
package ip_a_pkg {
  import addr_reg_pkg::*;
  struct trait_ip_a_s : addr_trait_s {
     rand int in [0..1] id;
  }
  component ip_a_c {
    action write_a {
      rand addr_claim_s<trait_ip_a_s> claim;
      constraint claim.size in [8, 16];
      constraint claim.alignment == 4;
      rand bit[32] data;
      exec body {
        addr_handle_t handle;
        handle = make_handle_from_claim(claim);
        write32(handle, data);
      }
    }
  }
}


package ip_b_pkg {
  import addr_reg_pkg::*;
  struct trait_ip_b_s : addr_trait_s {
     rand int in [0..127] id;
  }
  component ip_b_c {
    action write_a {
      rand addr_claim_s<trait_ip_b_s> claim;
      constraint claim.size in [2];
      rand bit[32] data;
      exec body {
        addr_handle_t handle;
        handle = make_handle_from_claim(claim);
        write32(handle, data);
      }
    }
  }
}
```

Example 343—Address space group

Copyright © 2026 Accellera. All rights reserved.

```
component pss_top {
  …
  ip_a_c ip_a;
  ip_b_c ip_b;
  transparent_addr_space_c<trait_ip_a_s> mema;
  transparent_addr_space_c<trait_ip_b_s> memb;

  addr_space_group_c mem_group;

  exec init_up {

      transparent_addr_region_s<trait_ip_a_s> region_ip_a_0, region_ip_a_1;
      region_ip_a_0.size = 128;
      region_ip_a_0.addr = 0x0;
      region_ip_a_0.trait.id = 0;
      (void)mema.add_region(region_ip_a_0);

      region_ip_a_1.size = 128;
      region_ip_a_1.addr = 128;
      region_ip_a_1.trait.id = 1;
      (void)mema.add_region(region_ip_a_0);
      mem_group.add_addr_space(mema);


      transparent_addr_region_s<trait_ip_b_s> region_ip_b[128];
      repeat(i:128) {
        region_ip_b[i].addr = i*2;
    region_ip_b[i].size = 2;
    region_ip_b[i].trait.id = i;
        (void)memb.add_region(region_ip_b[i]);
      }
      mem_group.add_addr_space(memb);

  }
  action entry_a {
    activity {
      parallel {
        replicate (1)
          do ip_a_c::write_a;
        replicate (2)
          do ip_b_c::write_a;
      }
    }
  }
}
```

Example 343—Address space group (cont.)

21.13 Data layout and access operations

21.13.1 Data layout

Many PSS use cases require writing structured data from the PSS model to byte-addressable space in a well-defined layout. In PSS, structured data is represented with a **struct**. For example, a DMA engine might expect DMA descriptors that encapsulate DMA operation to be in memory in a known layout. *Packed structs* may be beneficial to represent bit fields of hardware registers.

Copyright © 2026 Accellera. All rights reserved.

The built-in PSS library **struct `packed_s`** is used as a base **struct** to denote that a PSS **struct** is *packed*.

Any struct derived from built-in struct **`packed_s`** directly or indirectly is considered packed by the PSS tool. Packed structs are only allowed to have fields of numeric types, Boolean types, enumerated types that have a base type, packed struct types, or arrays thereof. Following are the declarations of the endianness **enum** and packed **struct** in **`std_pkg`**:

```
enum endianness_e {LITTLE_ENDIAN, BIG_ENDIAN};

struct packed_s <endianness_e e = LITTLE_ENDIAN> {};
```

Syntax 145—Packed_s base struct

Type extensions of packed structs shall not add new fields.

21.13.1.1 Packing rule

PSS uses the de facto packing algorithm from the GNU C/C++ compiler. The ordering of fields of structs follows the rules of the C language. This means that fields declared first would go in lower addresses. For this purpose, if a packed struct is derived from another packed struct, fields declared in the derived struct are considered to be declared later than those declared in the base struct. The layout of fields in a packed struct is defined by the endianness template parameter of the packed struct. Bit fields in PSS structs can be of any size. For this purpose, Boolean fields are considered to be of 1 bit.

For the packing algorithm, a register of size N bytes is used, where N*8 is greater than or equal to the number of bits in the packed struct.

For big-endian mode, fields are packed into registers from the most significant bit (MSB) to the least significant bit (LSB) in the order in which they are defined. Fields are packed in memory from the most significant byte (MSbyte) to the least significant byte (LSbyte) of the packed register. If the total size of the packed struct is not an integer multiple of bytes, don’t-care bits are added at the LSB side of the packed register.

For little-endian mode, fields are packed into registers from the LSB to the MSB in the order in which they are defined and packed in memory from the LSbyte to the MSbyte of the packed register. If the total size of the packed struct is not an integer multiple of bytes, don’t-care bits are added at the MSB side of the packed register.

21.13.1.2 Little-endian packing example

A packed struct is shown in Example 344. This struct has 30 bits. A register for packing this struct would have 4 bytes.

```
struct my_packed_struct : packed_s<LITTLE_ENDIAN> {
   bit[6] A;
   bit[2] B;
   bit[9] C;
   bit[7] D;
   bit[6] E;
}
```

Example 344—Packed PSS little-endian struct

Copyright © 2026 Accellera. All rights reserved.

Register packing will start from field A. The least significant bit of A would go in the least significant bit of the register, as shown in Figure 54. Field B would go after field A. The least significant bit of B would go in the lowest bit after A in the packed register, and so on. The layout of the packed struct in byte-addressable space is shown in Figure 55. (X means "don't-care bit" in Figure 54 and Figure 55.)

```
MSB                                                                 LSB
 X X E E E E E E D D D D D D D C C C C C C C C C B B A A A A A A
 X X 5 4 3 2 1 0 6 5 4 3 2 1 0 8 7 6 5 4 3 2 1 0 1 0 5 4 3 2 1 0
```

Figure 54—Little-endian struct packing in register

```
     byte 0           byte 1           byte 2           byte 3
B B A A A A A A  C C C C C C C C  D D D D D D D C  X X E E E E E E
1 0 5 4 3 2 1 0  7 6 5 4 3 2 1 0  6 5 4 3 2 1 0 8  X X 5 4 3 2 1 0
```

Figure 55—Little-endian struct packing in byte-addressable space

21.13.1.3 Big-endian packing example

A packed struct is shown in Example 345. This struct has 30 bits. A register for packing this struct would have 4 bytes.

```
struct my_packed_struct : packed_s<BIG_ENDIAN> {
   bit[6] A;
   bit[2] B;
   bit[9] C;
   bit[7] D;
   bit[6] E;
}
```

Example 345—Packed PSS big-endian struct

Register packing will start from field A. The most significant bit of A would go in the most significant bit of the register, as shown in Figure 56. Field B would go after field A. The most significant bit of B would go in the highest bit after A in the packed register, and so on. The layout of the packed struct in byte-addressable space is shown in Figure 57. (X means "don't-care bit" in Figure 56 and Figure 57.)

```
MSB                                                                 LSB
 A A A A A A B B C C C C C C C C C D D D D D D D E E E E E E X X
 5 4 3 2 1 0 1 0 8 7 6 5 4 3 2 1 0 6 5 4 3 2 1 0 5 4 3 2 1 0 X X
```

Figure 56—Big-endian struct packing in register

```
     byte 0           byte 1           byte 2           byte 3
A A A A A A B B  C C C C C C C C  C D D D D D D D  E E E E E E X X
5 4 3 2 1 0 1 0  8 7 6 5 4 3 2 1  0 6 5 4 3 2 1 0  5 4 3 2 1 0 X X
```

Copyright © 2026 Accellera. All rights reserved.

Figure 57—Big-endian struct packing in byte-addressable space

21.13.2 Sizeof_s

The template struct **sizeof_s** is used to query the physical storage size of a PSS data type. It applies to types that can be written to or read from a byte-addressable address space, namely numeric types, Booleans, enumerated types that have a base type, packed structs, and arrays thereof. The **sizeof_s** struct is declared in the **std_pkg** package.

21.13.2.1 Definition

```
struct sizeof_s<type T> {
  static const int nbytes = /* implementation-specific */;
  static const int nbits = /* implementation-specific */;
};
```

Syntax 146—Sizeof_s struct

The static constant **nbytes** is initialized to the number of consecutive addresses required to store a value of type **T** in a byte-addressable address space. When using the read/write target functions (see 21.13.9), this number of bytes is assumed to be taken up by the data in the target storage. For types that are not byte-aligned in size, the number of bytes is rounded up. For the definition of packed struct layout in an address space, see 21.13.1.

The static constant **nbits** is initialized to the exact number of bits that are taken up by the representation of a value of type **T** in a byte-addressable address space.

sizeof_s<> shall not be parameterized with types other than numeric types, Booleans, enumerated types that have a base type, packed structs, and arrays thereof.

21.13.2.2 Examples

The following code snippets show the value of **nbytes** of **sizeof_s<>** instantiated for several different types:

```
sizeof_s<int>::nbytes == 4

sizeof_s<int[4]>::nbytes == 1

sizeof_s<bit>::nbytes == 1
sizeof_s<bit[33]>::nbytes == 5

sizeof_s<array<int,10>>::nbytes == 40

struct my_packed_s : packed_s<> {bit[2] kind; int data;};
sizeof_s<my_packed_s>::nbytes  == 5
```

21.13.3 Address space handles

The built-in package **addr_reg_pkg** defines PSS types for *address space handles*.

Copyright © 2026 Accellera. All rights reserved.

```
typedef chandle addr_handle_t;

const addr_handle_t nullhandle = /* implementation-specific */;

struct sized_addr_handle_s <int SZ, // in bits
                            int lsb = 0,
                            endianness_e e = LITTLE_ENDIAN
                           > : packed_s<e> {
  addr_handle_t hndl;
};
```

Syntax 147—Address space handle

21.13.3.1 Generic address space handle

`addr_handle_t` is the generic type for address handles within an address space. A variable of type **`addr_handle_t`** resolves to a concrete address value during test execution, on the target platform. However, the concrete value of an address handle cannot be obtained during the solve process, on the solve platform. A field of type **`addr_handle_t`** cannot be declared directly in a packed struct type. Packed structs are defined in 21.13.1.

21.13.3.2 Nullhandle

`nullhandle` represents the address value `0` within the target address space, regardless of the actual mapping of regions.

21.13.3.3 Sized address space handle

The wrapper struct **`sized_addr_handle_s`** is used for specifying the size of an address handle in a packed struct. An address field within a packed struct shall only be declared using **`sized_addr_handle_s`**, and not directly as a field of type **`addr_handle_t`**.

The **`SZ`** parameter specifies the size of the handle itself in bits when used in a packed struct. Note that the **`SZ`** parameter is not the size of the data it is pointing to.

The **`lsb`** parameter defines the starting bit in the resolved address that would become bit `0` of sized address handle in packed struct. For example, assume that the resolved address is 64 bits and the size of the handle is 30 bits, with the **`lsb`** parameter set to 2. In this case, a sized handle in a packed struct would have bits 31 to 2 from the resolved address.

See an example in 21.13.10.

21.13.4 Obtaining an address space handle

A handle in an address space can be created from an address claim (with an optional offset value), from another handle (with an offset value), or from a region in an address space. An address claim is made using a claim struct declaration in actions and objects.

Copyright © 2026 Accellera. All rights reserved.

Some address space regions are non-allocatable. These regions can be used to represent memory-mapped I/O (MMIO) register spaces. A handle can be created from a region in an address space, in order to access non-allocatable regions.

A handle to a region is obtained when the region is added to the address space, using the **`add_region`** (see 21.10.1.2.1) or **`add_nonallocatable_region`** (see 21.10.1.2.2) functions. To create address handles from address claims or from other handles, the following functions are defined in the built-in package **`addr_reg_pkg`**.

21.13.4.1 Make_handle_from_claim function

The function **`make_handle_from_claim()`** creates an address handle from a claim, with an optional offset value.

```
function addr_handle_t make_handle_from_claim
  (addr_claim_base_s claim, bit[64] offset = 0, bool sub = false);
```

Syntax 148—Make_handle_from_claim function

The **`make_handle_from_claim`** function arguments are:

- A claim struct instance declared in an action or a flow/resource object
- An optional offset value, of a 64-bit type
- An optional offset subtraction choice (sub), of type bool

The returned handle's resolved address will be the sum of the claim's resolved address and the offset if sub is false (default) and the claim's resolved address minus the offset if sub is true. The return value of the function is of type **`addr_handle_t`**.

Example 346 shows the usage for the `make_handle_from_claim` function.

```
action my_action {
  rand transparent_addr_claim_s<> claim;

  constraint claim.size == 128;
  constraint claim.alignment == 2**4;

  exec body {
    int offset = 16;
    int data = 128;

    addr_handle_t h0 = make_handle_from_claim(claim);
    write32(h0, data); // access API defined in 21.13.9.1

    // Address handle from claim with an offset
    addr_handle_t h1 = make_handle_from_claim(claim, offset);
    write32(h1, data);
  }
};
```

Example 346—Make_handle_from_claim example

Copyright © 2026 Accellera. All rights reserved.

21.13.4.2 Make_handle_from_handle function

The function **`make_handle_from_handle()`** creates an address handle from another handle, given an offset.

```
function addr_handle_t make_handle_from_handle
  (addr_handle_t handle, bit[64] offset, bool sub = false);
```

Syntax 149—Make_handle_from_handle function

The **`make_handle_from_handle`** function arguments are:

- A handle that was created by a different call to a **`make_handle`** function
- An offset value, of a 64-bit type
- An optional offset subtraction choice (sub), of type bool

The returned handle's resolved address will be the sum of the **`handle`** parameter's resolved address and the offset if sub is false (default) and the handle parameter's resolved address minus the offset if sub is true. The return value of the function is of type **`addr_handle_t`**.

Example 347 shows how usage of a negative offset can assist in creating different virtual addresses for different cpu's using the same physical address space. When accessing region1, the `cpu_core[0]` offsets will be positive, and the `cpu_core[0]` offset will be negative. The offsets and their direction are configured at the top level in `pss_top`.

Copyright © 2026 Accellera. All rights reserved.

```
component cpu_c {
     int myid;
     bit[32] region_offset;
     bool is_region_offset_negative;
     action cpu_config_a {
        lock executor_r executor_1;
        rand bit[64] in [0, 256] data;
        rand transparent_addr_claim_s<trait_s> claim;
        constraint claim.size in [8, 16];
        constraint claim.alignment == 4;

        addr_handle_t handle;
        exec pre_body {
          handle = make_handle_from_claim(claim);
        // Region 1 is virtualized differently for each cpu instance
        // differently
        if (get_tag(handle) == "region1") {
          handle = make_handle_from_handle(handle, comp.region_offset,
                     comp.is_region_offset_negative);
        }
       }
       exec body {
          addr_reg_pkg::write64(handle, data);
       }
     }
}
component pss_top {
  cpu_c cpu_core[NOF_EXECUTORS];
  transparent_addr_space_c<trait_s> mem;
  exec init_down {

    foreach(cpu_entry:cpu_core[i]) {
      cpu_entry.myid = i;
    }
    cpu_core[0].region_offset = 0x2000;
    cpu_core[0].is_region_offset_negative = false;
    cpu_core[1].region_offset = 0x2000;
    cpu_core[1].is_region_offset_negative = true;

    transparent_addr_region_s<trait_s> region0, region1;
    region0.size = 128;
    region0.addr = 0x10000;
    region0.trait.id = 0;
    region0.tag = "region0";
    (void)mem.add_region(region0);
    region1.size = 128;
    region1.trait.id = 1;
    region1.addr = 0x20000;
    region1.tag = "region1";
    (void)mem.add_region(region1);
}
```

Example 347—Negative offset usage

Copyright © 2026 Accellera. All rights reserved.

```
action entry_a {
  activity {
    foreach (comp.cpu_core[i])
      do cpu_c::cpu_config_a with {claim.trait.id == i; executor_1.trait.id
   == i;};
       }
     }
   }
```

Example 347—Negative offset usage (cont.)

Example 348 shows the usage for the `make_handle_from_handle` function.

```
action my_action {
   rand transparent_addr_claim_s<> claim;
   constraint claim.alignment == 2**4;

   exec body {
      int offset = 16;
      int data = 128;

      addr_handle_t h0 = make_handle_from_claim(claim, offset);
      write32(h0, data);

      // Make handle from another handle with an offset
      addr_handle_t h1 = make_handle_from_handle(h0, sizeof_s<int>::nbytes);
      write32(h1, data);
   }
};
```

Example 348—Make_handle_from_handle example

21.13.5 Addr_value function

The function **`addr_value()`** returns the resolved address of the parameter handle, as a numeric value. **`addr_value()`** is a target function and shall only be used in **exec body**, **run_start**, **run_end**, or functions called from these **exec** blocks.

```
target function bit[64] addr_value(addr_handle_t hndl,
                                   const mem_access_desc_s mem_access_desc = {});
```

Syntax 150—Addr_value function

Per-executor custom implementations of the **`addr_value()`** function may be provided, much as custom implementations of read/write functions are (see 21.13.9.5).

21.13.6 Addr_value_solve function

```
solve function bit[64] addr_value_solve(addr_handle_t hndl,
                                        const mem_access_desc_s mem_access_desc = {});
```

Syntax 151—Addr_value_solve function

Copyright © 2026 Accellera. All rights reserved.

The solve function **addr_value_solve()** returns either the full absolute address of the **hndl** parameter or the offset of the **hndl** parameter within its containing address region as a numeric value. If the **hndl** parameter is within a transparent region, the returned value will be an absolute address. If the **hndl** parameter is within an opaque region, the returned value *may* be an absolute address or an offset depending on what tool-specific metadata has been supplied to the PSS processing tool. The **addr_value_abs()** function is used to determine what information will be returned by **addr_value_solve()** for a given address handle.

Users may provide executor-specific implementations of **addr_value_solve()** by overriding this method in an executor implementation.

The **addr_value_solve()** function may only be called in the context of a **pre_body** *exec block*. If **addr_value_solve()** is called from other contexts, the return value is undefined.

21.13.7 Addr_value_abs function

```
solve function bool addr_value_abs(addr_handle_t hndl,
                                   const mem_access_desc_s mem_access_desc = {});
```

Syntax 152—Addr_value_abs function

The solve function **addr_value_abs()** returns 'true' if the absolute address value is available for the specified address handle. The absolute address value is available if **hndl** is within a transparent region, and *may* be available when **hndl** is within an opaque region depending on what tool-specific metadata has been supplied to the PSS processing tool.

Users may provide executor-specific implementations of **addr_value_abs()** by overriding this method in an executor implementation.

The **addr_value_abs()** function may only be called in the context of a **pre_body** *exec block*. If **addr_value_abs()** is called from other contexts, the return value is undefined.

21.13.8 Get_tag function

The function **get_tag()** returns the *tag* (see Syntax 153) of the region in which the specified address handle is located. **get_tag()** shall only be used in **exec pre_body**, **body**, **run_start**, **run_end**, or in functions called from these **exec** blocks.

```
function string get_tag(addr_handle_t hndl);
```

Syntax 153—Get_tag function

21.13.9 Access operations

Read/write operations of PSS data from/to byte-addressable address space are defined as a set of target functions. Target **exec** blocks (**exec body**, **run_start**, **run_end**), and functions called from them, may call these core library functions to access allocated addresses.

Access functions use an address handle to designate the required location within an address space (see Syntax 154 and refer to Annex C.3).

Copyright © 2026 Accellera. All rights reserved.

```
struct mem_access_desc_s {}
target function bit[8] read8(addr_handle_t hndl,
                             const mem_access_desc_s mem_access_desc = {});
target function bit[16] read16(addr_handle_t hndl,
                               const mem_access_desc_s mem_access_desc = {});
target function bit[32] read32(addr_handle_t hndl,
                               const mem_access_desc_s mem_access_desc = {});
target function bit[64] read64(addr_handle_t hndl,
                               const mem_access_desc_s mem_access_desc = {});
```

Syntax 154—Mem_access_desc_s

PSS provides a way for a tool or user to customize the implementation of access functions based on user input. The access functions have an optional parameter of type `mem_access_desc_s` declared in the `addr_reg_pkg` (see Annex C.3). This argument is used by native PSS implementations of access functions (see 21.13.9.5). Tools may ignore this argument when directly mapping the `addr_reg` package memory access functions to a target language. The signature of the `addr_reg` package memory access functions has the third argument to enable the global calling of the memory access functions from executors mentioned (see 21.13.9.5).

The tool or user can extend the core library struct `mem_access_desc_s {}` to provide additional controls to the access functions.

Example 349 shows how a PSS user can use the additional argument to customize whether a `write64` should be done in a single 64-bit write or two 32-bit writes.

Copyright © 2026 Accellera. All rights reserved.

```
import addr_reg_pkg::*;
extend struct mem_access_desc_s {
  rand bool single_write;
}
component c_executor_c : executor_base_c {
  target function void write64(addr_handle_t hndl, bit[64] data,
                               const mem_access_desc_s mem_access = {}) {
    if (mem_access.single_write)
      addr_reg_pkg::write64(hndl, data);
    else {
      addr_reg_pkg::write32(make_handle_from_handle(hndl, 4), data[63:32]);
      addr_reg_pkg::write32(hndl, data[31:0]);
    }
  }
}
component pss_top {
  action write_a {
    rand mem_access_desc_s mem_access;
    rand transparent_addr_claim_s<trait_s> claim;
    constraint claim.size in [8, 16];
    constraint claim.alignment == 4;
    constraint mem_access.single_write == false;
    rand bit[64] data;
    addr_handle_t hndl;
    exec pre_body {
      hndl = make_handle_from_claim(claim);
    }
  }
  exec body {
    write64(hndl, data, mem_access);
    write64(hndl, data, mem_access);
  }
}
```

Example 349—Customized access functions

Example 350 shows how a PSS user can use the additional argument to customize whether a write64 should be written frontdoor (through a verification agent) or written backdoor (directly to memory).

Copyright © 2026 Accellera. All rights reserved.

```
import addr_reg_pkg::*;
extend struct mem_access_desc_s {
  rand bool frontdoor;
  constraint default frontdoor == true;
}
import target function void tb_hdl_deposit(string path, bit[64] index,
                                           bit[32] data);
component sv_executor_c : executor_base_c {
  target function void write64(addr_handle_t hndl, bit[64] data,
                               const mem_access_desc_s mem_access = {}) {
    if (mem_access.frontdoor)
      addr_reg_pkg::write64(hndl, data);
    else {
      tb_hdl_deposit("top.system.memory_subsystem.mem", addr_value(hndl),
                     data);
  }
}

component pss_top {
  action write_a {
    rand mem_access_desc_s frontdoor_access;
    rand transparent_addr_claim_s<trait_s> claim;
    constraint claim.size in [8, 16];
    constraint claim.alignment == 4;
    constraint default disable frontdoor_access.frontdoor ;
    constraint frontdoor_access.frontdoor == false;
    rand bit[64] data;
    addr_handle_t hndl;
    exec pre_body {
      hndl = make_handle_from_claim(claim);
    }
    exec body {
      write64(hndl, data, frontdoor_access);
    }
  }
}
```

Example 350—Customized access functions (frontdoor or backdoor)

21.13.9.1 Primitive read operations

Syntax 155 defines read operations for integer types from byte addressable address spaces to read one, two, four or eight consecutive bytes starting at the address indicated by the **`addr_handle_t`** argument.

Copyright © 2026 Accellera. All rights reserved.

```
target function bit[8] read8(addr_handle_t hndl,
                             const mem_access_desc_s mem_access_desc = {});
target function bit[16] read16(addr_handle_t hndl,
                               const mem_access_desc_s mem_access_desc = {});
target function bit[32] read32(addr_handle_t hndl,
                               const mem_access_desc_s mem_access_desc = {});
target function bit[64] read64(addr_handle_t hndl,
                               const mem_access_desc_s mem_access_desc = {});
```

Syntax 155—Primitive read operations for byte addressable spaces

The first byte goes into bits [7:0], then the next byte goes into bits [15:8], and so on.

21.13.9.2 Primitive write operations

Syntax 156 defines write operations for integer types to byte addressable address spaces to write one, two, four or eight consecutive bytes from the **`data`** argument starting at the address indicated by the **`addr_handle_t`** argument.

```
target function void write8 (addr_handle_t hndl, bit[8] data,
                             const mem_access_desc_s mem_access_desc = {});
target function void write16(addr_handle_t hndl, bit[16] data,
                             const mem_access_desc_s mem_access_desc = {});
target function void write32(addr_handle_t hndl, bit[32] data,
                             const mem_access_desc_s mem_access_desc = {});
target function void write64(addr_handle_t hndl, bit[64] data,
                             const mem_access_desc_s mem_access_desc = {});
```

Syntax 156—Primitive write operations for byte addressable spaces

Bits [7:0] of the input **`data`** go into the starting address specified by the **`addr_handle_t`** argument, bits [15:8] go into the next address (starting address + 1), and so on.

21.13.9.3 Read and write N consecutive bytes

Syntax 157 defines operations to read and write a series of consecutive bytes from byte addressable space.

For a read operation, the read data is stored in the argument **`data`**. For function **`read_bytes()`**, the **`size`** argument indicates the number of consecutive bytes to read. The returned list is resized accordingly, and its previous values, if any, are overwritten.

For a write operation, the input data is taken from the argument **`data`**. For function **`write_bytes()`**, the number of bytes to write is determined by the list size of the **`data`** parameter.

Copyright © 2026 Accellera. All rights reserved.

```
target function void read_bytes (addr_handle_t hndl, list<bit[8]> data,
                                 int size, const mem_access_desc_s mem_access_desc = {});
target function void write_bytes(addr_handle_t hndl, list<bit[8]> data,
                                 const mem_access_desc_s mem_access_desc = {});
```

Syntax 157—Read and write series of bytes

The first byte read comes from the address indicated by the **hndl** argument. This byte is stored at the first location (index 0) in the **data** list. The second byte comes from the address incremented by one and is stored at the second location (index 1) in the **data** list, and so on. The same semantics apply to **write_bytes**().

21.13.9.4 Read and write packed structs

Read and write operations to access packed structs are defined in Syntax 158. Argument **packed_struct** of functions **read_struct()** and **write_struct()** shall be a subtype of the **packed_s** struct. The **packed_struct** argument is read from or written to the address specified by the **hndl** argument.

```
target function void read_struct(addr_handle_t hndl, struct packed_struct,
                                 const mem_access_desc_s   mem_access_desc = {});
target function void write_struct(addr_handle_t hndl, struct packed_struct,
                                 const mem_access_desc_s   mem_access_desc = {});
```

Syntax 158—Read and write packed structs

The PSS implementation shall convert calls to **read_struct()** and **write_struct()** to one or more invocations of the primitive read and write operations (see 21.13.9.1 and 21.13.9.2) or to an invocation of the **read_bytes()** or **write_bytes()** function (see 21.13.9.3). Reading and writing of structs of size 8, 16, 32, or 64 bits stored at a correspondingly aligned address shall be implemented with a single primitive operation of the corresponding size, and in other cases may be partitioned into one or more primitive operations of any size, or a single call to the **read_bytes()** or **write_bytes()** function.

The **mem_access_desc_s** argument will be passed as is to the used primitive functions.

21.13.9.5 Executor-based customization of memory functions

PSS tools may provide built-in implementations of read, write, and address resolution operations for mainstream execution contexts. However, users can optionally customize the implementation of these operations for their own purposes and execution contexts.

Calls to primitive read, write, and address resolution functions (defined above in 21.13.9.1, 21.13.9.2, and 21.13.5), and calls to byte list read/write functions (defined above in 21.13.9.3), are delegated to functions with the identical prototype in the executor instance assigned to the evaluation action or flow/resource object. When a read or write memory access function is invoked, the computation of the effective address is delegated to **addr_value()**. If the user provides an override of **addr_value()**, the overridden implementation shall be used. Syntax 159 below shows the declarations of the executor implementation functions.

Copyright © 2026 Accellera. All rights reserved.

```
extend component executor_base_c {
  target function bit[64] addr_value(addr_handle_t hndl,
                                     const mem_access_desc_s mem_access_desc = {});
  solve function bit[64] addr_value_solve(addr_handle_t hndl,
                                          const mem_access_desc_s mem_access_desc = {});
  solve function bool addr_value_abs(addr_handle_t hndl,
                                     const mem_access_desc_s mem_access_desc = {});
  target function bit[8] read8 (addr_handle_t hndl,
                                const mem_access_desc_s mem_access_desc = {});
  target function bit[16] read16(addr_handle_t hndl,
                                 const mem_access_desc_s mem_access_desc = {});
  target function bit[32] read32(addr_handle_t hndl,
                                 const mem_access_desc_s mem_access_desc = {});
  target function bit[64] read64(addr_handle_t hndl,
                                 const mem_access_desc_s mem_access_desc = {});
  target function void write8(addr_handle_t hndl, bit[8] data,
                              const mem_access_desc_s mem_access_desc = {});
  target function void write16(addr_handle_t hndl, bit[16] data,
                               const mem_access_desc_s mem_access_desc = {});
  target function void write32(addr_handle_t hndl, bit[32] data,
                               const mem_access_desc_s mem_access_desc = {});
  target function void write64(addr_handle_t hndl, bit[64] data,
                               const mem_access_desc_s mem_access_desc = {});
  target function void read_bytes (addr_handle_t hndl, list<bit[8]> data, int size,
                                   const mem_access_desc_s mem_access_desc = {});
  target function void write_bytes(addr_handle_t hndl, list<bit[8]> data,
                                   const mem_access_desc_s mem_access_desc = {});
};
```

Syntax 159—Primitive operation implementation functions

Note that struct read/write functions (defined above in 21.13.9.4) and register read/write functions (defined below in 21.14.1) are implemented in terms of their respective primitive operations. Therefore, custom implementations of the primitive operations in an executor apply similarly to struct and register read/write functions.

The code in Example 351 below illustrates how a PSS implementation may define the delegation of one of the primitive read/write functions to the corresponding function in the current executor. The actual implementation does not necessarily take this form, but should have equivalent observable behavior. See 21.7.2.5 for more on the semantics of function **`executor()`**.

Copyright © 2026 Accellera. All rights reserved.

```
function bit[32] read32(addr_handle_t hndl) {
  if (executor() != null ) {
     return executor().read32(hndl);
  } else {
     // return value per default implementation
  }
}
```

Example 351—Illustration of read32()

Example 352 below demonstrates how primitive operations **`read32()`** and **`write32()`** are mapped to calls to functions of a C bus transactor in the context of a user-defined executor type.

```
function bit[32] my_transactor_read_word(bit[64] addr);
import target C function my_transactor_read_word;

function void my_transactor_write_word(bit[64] addr, bit[32] data);
import target C function my_transactor_write_word;

component my_transactor_executor_c
    <struct TRAIT : executor_trait_s = empty_executor_trait_s>
    : executor_c<TRAIT> {
  function bit[32] read32(addr_handle_t hndl,
                          const mem_access_desc_s mem_access_desc = {}) {
    return my_transactor_read_word(addr_value(hndl));
  }

  function void write32(addr_handle_t hndl, bit[32] data,
                        const mem_access_desc_s mem_access_desc = {}) {
    my_transactor_write_word(addr_value(hndl), data);
  }
};
```

Example 352—Mapping of primitive operations to foreign C functions

In Example 353 below, executor type `uvm_ubus_executor_c` corresponds to a UVM bus master. The **`write8()`** function is defined in terms of a SystemVerilog imported function (task) that starts a write-byte sequence on the agent designated by the path parameter. The executor type is instantiated twice under **`pss_top`**, and each instance is associated with a different UVM agent in the target environment using the UVM path.

Copyright © 2026 Accellera. All rights reserved.

```
import target SV function void ubus_write8(string uvm_path, bit[64] addr,
   bit[8] data);

component uvm_ubus_executor_c : executor_c<bus_trait_s> {
   string uvm_path;

   function void write8(addr_handle_t hndl, bit[8] data,
                        const mem_access_desc_s mem_access_desc = {}) {
      ubus_write8(uvm_path, addr_value(hndl), data);
   }
};

extend component pss_top {
  uvm_ubus_executor_c masters[2];
  executor_group_c<bus_trait_s> bus_group;
  exec init_down {
    foreach (m: masters) {
      bus_group.add_executor(m);
    }
    masters[0].uvm_path = "uvm_test_top.env.ubus_master0";
    masters[1].uvm_path = "uvm_test_top.env.ubus_master1";
  }
};
```

Example 353—Mapping of primitive operations to UVM sequences

In Example 354 below, an executor corresponding to a 32-bit architecture CPU customizes the **`read64()`** and **`write64()`** operations to be implemented in terms of the built-in **`read32()`** and **`write32()`** operations.

```
component my_32bit_cpu_c : executor_c<my_core_trait_s> {
  function bit[64] read64(addr_handle_t hndl) {
    bit[64] result;
    result[31: 0] = read32(hndl);
    result[63:32] = read32(make_handle_from_handle(hndl,4));
    return result;
   }

   function void write64(addr_handle_t hndl, bit[64] data,
                         const mem_access_desc_s mem_access_desc = {}) {
      write32(hndl, data[31:0]);
      write32(make_handle_from_handle(hndl,4), data[63:32]);
   }
};
```

Example 354—Implementing primitive operations in terms of other operations

In the example below, the user has an address map where each of a set of executors is allocated a unique set of addresses within the address space. While each executor is assigned a unique portion of the global address space, the executor-specific address window is mapped at the same address from the perspective of the executor. Allocations are modeled using the global address map to ensure claims are globally unique. However, depending on the executor, an address may need to be transformed to conform to the executor-specific address map.

Copyright © 2026 Accellera. All rights reserved.

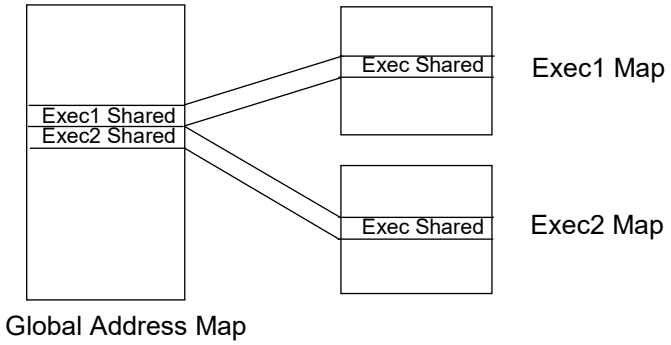


Figure 58—Executor address mapping

Overriding the **addr_value()** function can be used to perform such custom translations. The **addr_window_exec_c** executor shown below overrides the **addr_value()** function and applies a translation if the address falls within a specific window that is configurable on a per-executor instance basis.

Let's assume that the executor-specific address windows are located at `0x80000000` and `0x80001000` in the global address map. Each executor maps this shared window at `0x1000`. The executor instantiation and configuration below show how we could configure this translation scheme. When, for example, an action running on `exec1` accesses address `0x8000_0100`, the customized **addr_value()** function will convert the address to `0x0000_1100`.

```
component addr_window_exec_c : executor_c<> {
  bit[64] window_base   = 0x80000000;
  bit[64] window_size   = 0x1000;
  bit[64] window_offset = 0x80000000;

  function bit[64] addr_value(addr_handle_t hndl,
                              const mem_access_desc_s mem_access_desc = {}) {
    bit[64] addr = super.addr_value(hndl);
    if (addr >= window_base && addr < (window_base+window_size)) {
        addr = (addr-window_offset)+0x1000;
    }
    return addr;
  }
}

component subsystem_c {
  addr_window_exec_c exec1;
  addr_window_exec_c exec2;

  exec init_down {
    exec1.window_base   = 0x8000_0000;
    exec1.window_offset = 0x8000_0000;
    exec2.window_base   = 0x8000_1000;
    exec2.window_offset = 0x8000_1000;
  }
}
```

Example 355—Customization of addr_value()

Copyright © 2026 Accellera. All rights reserved.

Example 356 shows overriding **addr_value()** to perform address translation using attributes in **mem_access_desc_s**. The executor calls the base implementation to obtain a physical address, and conditionally translates it based on the privilege field. The PSS tool propagates **mem_access_desc_s** from **write32()** to **addr_value()**, and uses the returned address for the access, avoiding the need to override individual read/write functions.

```
import addr_reg_pkg::*;
enum privilege_level_e {MACHINE, SUPERVISOR};

extend struct mem_access_desc_s {
  rand privilege_level_e privilege;
};

import target bit[64] function supervisor_physical_to_virtual_addr(bit[64]
                                                                   physical_addr);

component my_executor_c : executor_base_c {
  target function bit[64] addr_value(addr_handle_t hndl,
                                     const mem_access_desc_s mem_access_desc = {}) {
    bit[64] physical_addr = super.addr_value(hndl, mem_access_desc);
    if (mem_access_desc.privilege == MACHINE)
      return physical_addr;
    else
      return supervisor_physical_to_virtual_addr(physical_addr);
    }
  }
}
component pss_top {
  action write_a {
    rand bit[32] data;
    rand transparent_addr_claim_s<> claim;
    exec body {
      mem_access_desc_s desc;
      desc.privilege = SUPERVISOR;
      addr_handle_t h0 = make_handle_from_claim(claim);
      write32(h0, data, desc);
    }
  }
}
```

Example 356—Customization of memory functions

Copyright © 2026 Accellera. All rights reserved.

21.13.10 Target data structure setup example

The following example demonstrates use of packed PSS data written to allocations on byte addressable space. It also demonstrates the use of address handles to construct complex data structures in target memory. Lifetime of allocation is extended by using address handles in flow objects.

```
buffer data_buff {
   rand addr_claim_s<> mem_seg;
};

component dma_c {

   struct descriptor_s : packed_s<> {
      sized_addr_handle_s<32> src_addr;
      sized_addr_handle_s<32> dst_addr;
      int size;
      sized_addr_handle_s<32> next_descr;
   };

   state descr_chain_state {
      list<addr_handle_t> handle_list;
   };

   pool descr_chain_state descr_chain_statevar;
   bind descr_chain_statevar *;

   action alloc_first_descr {
      output descr_chain_state out_chain;

      rand addr_claim_s<> next_descr_mem;
      constraint next_descr_mem.size == sizeof_s<descriptor_s>::nbytes;

      exec post_solve {
         out_chain.handle_list.push_back(
                                    make_handle_from_claim(next_descr_mem));
      }
   };
```

Example 357—Example using complex data structures

Copyright © 2026 Accellera. All rights reserved.

```
action chained_xfer {
   input  data_buff src_buff;
   output data_buff dst_buff;
   constraint dst_buff.mem_seg.size == src_buff.mem_seg.size;

   input  descr_chain_state in_chain;
   output descr_chain_state out_chain;

   rand bool last;

   descriptor_s descr;

   rand addr_claim_s<> next_descr_mem;
   constraint next_descr_mem.size == sizeof_s<descriptor_s>::nbytes;

   addr_handle_t descr_hndl;

   exec post_solve {
      descr.src_addr.hndl = make_handle_from_claim(src_buff.mem_seg);
      descr.dst_addr.hndl = make_handle_from_claim(dst_buff.mem_seg);
      descr.size = src_buff.mem_seg.size;
      if (last) {
         descr.next_descr.hndl = nullhandle;
      } else {
         descr.next_descr.hndl = make_handle_from_claim(next_descr_mem);
      }

      // tail of current list
      descr_hndl = in_chain.handle_list[in_chain.handle_list.size()-1];

      // copy over list from input to output
      out_chain.handle_list = in_chain.handle_list;
      // add next pointer
      out_chain.handle_list.push_back(
                              make_handle_from_claim(next_descr_mem));
   }

   exec body {
      write_struct(descr_hndl,descr);
   }
};

action execute_xfer {
   input descr_chain_state in_chain;

   addr_handle_t descr_list_head;

   exec post_solve {
      descr_list_head = in_chain.handle_list[0]; // head of list
   }

   exec body {
      // Initiate chained-transfer with descr_list_head
      // Wait for the chained-transfer to complete
   }
};
```

Example 357—Example using complex data structures (cont.)

Copyright © 2026 Accellera. All rights reserved.

```
    action multi_xfer {
       rand int in [1..10] num_of_xfers;

       activity {
          do alloc_first_descr;
          repeat (i: num_of_xfers) {
             do chained_xfer with {last == (i == num_of_xfers-1);};
          }
          do execute_xfer;
       }
    };
};
```

Example 357—Example using complex data structures (cont.)

In this example, the `chained_xfer` **action** represents the data flow (source/destination buffers) associated with this transaction. It populates the descriptor, including a pointer to the next descriptor, which it allocates. Its runtime execution writes the full descriptor out to memory, in the location allocated for it by the previous link in the chain.

21.14 Registers

A PSS model will often specify interaction with the hardware SUT to control how the PSS tool-generated code will read/write to programmable registers of the SUT. This section shows how to associate meaningful identifiers with register addresses that need to be specified in the PSS model description, as well as manipulation of the value of register fields by name.

All the core library constructs in this section are declared in the **`addr_reg_pkg`** package. For brevity, the definitions below do not include the package name.

21.14.1 PSS register definition

A *register* is a logical aggregation of fields that are addressed as a single unit.

The **`reg_c`** component is a base type for specifying the programmable registers of the DUT. Note that it is a **pure** component (see 9.1.7). It shall be illegal to extend the **`reg_c`** component and its base components.

```
enum reg_access {READWRITE, READONLY, WRITEONLY};
pure component reg_c <type R,
  reg_access ACC = READWRITE,
  int SZ = (8*sizeof_s<R>::nbytes)> : reg_sized_c<SZ> {
  target function R read(const mem_access_desc_s mem_access_desc = {});
  target function void write(R r,
                             const mem_access_desc_s mem_access_desc = {});
  target function void write_masked(R mask, R val,
                                    const mem_access_desc_s mem_access_desc = {});
};
```

Syntax 160—PSS register definition

Copyright © 2026 Accellera. All rights reserved.

Component **reg_c** is parameterized by:

a) A type **R** for the value (referred to as the *register-value type*) that can be read/written from/to the register, which can be:
   1) A packed structure type (that represents the register structure)
   2) A bit-vector type (**bit[N]**)

b) Kind of access allowed to the register, which by default is **READWRITE**

c) Width of the register (**SZ**) in number of bits, which by default equals the size of the register-value type **R** (rounded up to a multiple of 8)

SZ, if specified by the user, shall be greater than or equal to the size of the register-value type **R**. If the size of the register-value type **R** is less than the width of the register, it will be equivalent to having **SZ – sizeof_s<R>::nbits** reserved bits at the end of the structure.

```
pure component reg_sized_c<int SZ> : reg_base_c {
  target function bit[SZ] read_val(const mem_access_desc_s mem_access_desc = {});
  target function void write_val(bit[SZ] r, const mem_access_desc_s mem_access_desc = {});
  target function void write_val_masked(bit[SZ] mask, bit[SZ] val,
                                        const mem_access_desc_s mem_access_desc = {});
  target function void write_field(string name, bit[SZ] val,
                                   const mem_access_desc_s mem_access_desc = {});
  target function void write_fields(list<string> names, list<bit[SZ]> vals,
                                    const mem_access_desc_s mem_access_desc = {});
}
```

Syntax 161—Reg_sized_c definition

The *reg_sized_c* component provides a generic mechanism for accessing registers of a given width by value. The register access functions described in Syntax 161 and Syntax 162 may be called from the test-realization layer of a PSS model. Being declared as target functions, these need to be called in an **exec body** context.

```
pure component reg_base_c {
  function addr_handle_t get_handle();
}
```

Syntax 162—Reg_base_c definition

The *reg_base_c* component provides a generic mechanism for accessing base properties of registers, such as the corresponding address handle. The `get_handle()` function returns an address handle corresponding to the register. See 21.13.4 for more details on obtaining an address handle to registers and register groups.

Copyright © 2026 Accellera. All rights reserved.

```
struct CR : packed_s<> {
  bit en;
  bit[11] pad;
  bit[4] mode;
  bit[16] coeff;
}
pure component dut_regs_c : reg_group_c {
  reg_c<CR> cr1;
  reg_c<bit[32]> cr2;
  reg_c<bit[32]> cr3;
}
target function zero_r32(list<ref reg_sized_c<32>> regs) {
  foreach (r : regs) {
    r.write_val(0);
  }
}
component dut_c {
  dut_regs_c regs;
  action zero_a {
    exec body {
      zero_r32({comp.regs.cr1, comp.regs.cr2, comp.regs.cr3});
    }
  }
}
```

Example 358—Using sized register base

Example 358 shows how *reg_sized_c* can be used to operate on registers of the same size in a generic manner.

The **read()** and **read_val()** functions return the value of the register in the DUT (the former returns an instance of register-value type and the latter returns a bit vector). The **write()** and **write_val()** functions update the value of a register in a DUT (the former accepting an instance of register-value type and the latter a bit vector). If the register-value type is a bit vector, then the functions **read()** and **read_val()** are equivalent, as are **write()** and **write_val()**.

The **write_masked()** and **write_val_masked()** methods cause the register to be read, a write value to be calculated from the current register value and the specified masked value, and the write value to be written back to the register. The effect is the following:

```
REG_VAL(new) = (REG_VAL(current) & ~mask) | (val & mask)
```

If dedicated read-modify-write instructions are available on a platform, a PSS processing tool may, but is not required to, implement these operations in terms of those instructions.

The **write_masked()** and **write_val_masked()** methods only differ in how the mask and value are specified. In the case of **write_val_masked()**, both are specified as numeric quantities. In the case of **write_masked()**, both are specified in terms of the register-value type used to define the register.

The **write_field()** and **write_fields()** methods specify read-write-modify operations on a register using named register fields. Note that these methods may only be used on registers specified in terms of a struct data type. The following restrictions apply to the field names specified to **write_field()** and **write_fields()**:

a) Only string literals may be used in specifying field names.

Copyright © 2026 Accellera. All rights reserved.

b) The names may only specify top-level fields, and may not specify dotted hierarchical references.

c) The field name may not refer to aggregate data type fields within the register.

d) The set of strings passed to **`write_fields()`** shall be unique.

```
struct CR : packed_s<> {
    bit          en;
    bit[11]      pad;
    bit[4]       mode;
    bit[16]      coeff;
}

pure component dut_regs_c : reg_group_c {
    reg_c<CR>    cr;
}

component dut_c {
    dut_regs_c         regs;

    action cfg_a {
        rand bit[4]  mode;
        rand bit[16] coeff;
        exec body {
            // Three equivalent ways to modify the 'mode' and 'coeff' fields
            comp.regs.cr.write_masked(
                {.mode=~0, .coeff=~0}, {.mode=mode, .coeff=coeff});
            comp.regs.cr.write_val_masked(
                0xFFFFF000, (coeff << 16) | (mode << 12));
            comp.regs.cr.write_fields({"mode", "coeff"}, {mode, coeff});
        }
    }

    action enable_a {
        exec body {
            // Two equivalent ways to set the 'en' bit
            comp.regs.cr.write_masked({.en=~0}, {.en=1});
            comp.regs.cr.write_field("en", 1);
        }
    }
}
```

Example 359—Read-modify-write operations

In Example 359, a register is defined in terms of a packed struct with three operational fields and a reserved unused region (`pad`). In the action `cfg_a`, three different ways are shown to ensure that the `mode` and `coeff` fields are set to specific values while leaving the `en` field unmodified:

a) Mask and value parameters are formulated using struct literal expressions and passed to the **`write_masked()`** method. Fields in the mask parameter are set to the negation of 0 (all bits set) in order to cause the value of the corresponding register bits to be set. Unspecified fields in the mask parameter take on the default value, which PSS specifies as 0 for integer data types.

b) Numeric mask and value parameters are computed using shift and composition operations and passed to the **`write_val_masked()`** method.

c) Lists of field names and field values are passed to the **`write_fields()`** method.

Copyright © 2026 Accellera. All rights reserved.

See 21.14.5 for a description of the implementation of these functions. It shall be an error to call a register read or read-modify-write function on a register object whose access is set to **WRITEONLY**. It shall be an error to call a register write or read-modify-write function on a register object whose access is set to **READONLY**.

A template instantiation of the class **reg_c** (i.e., **reg_c<R, ACC, SZ>** for some concrete values for **R**, **ACC** and **SZ**) or a component derived from such a template instantiation (directly or indirectly) is a *register type*. An object of register type can be instantiated only in a *register group* (see 21.14.2).

Example 360 shows examples of register declarations.

```
struct my_reg0_s : packed_s<> { // (1)
   bit [16] fld0;
   bit [16] fld1;
};

pure component my_reg0_c : reg_c<my_reg0_s> {} // (2)

struct my_reg1_s : packed_s<> {

   bit     fld0;
   bit [2] fld1;
   bit [2] fld2[5];                // (3)
};

pure component my_reg1_c : reg_c<my_reg1_s, READWRITE, 32> {} // (4)
```

Example 360—Examples of register declarations

Notes:

1) `my_reg0_s` is the register-value type. The endianness can be explicitly specified if needed.
2) `my_reg0_c` is the register type. Since it derives from **reg_c**<`my_reg0_s`>, it inherits the **reg_c** read/write functions. Note that the access is **READWRITE** by default and the width equals the size of the associated register-value type, `my_reg0_s`.
3) Fixed-size arrays are allowed.
4) **sizeof_s<**`my_reg1_s`**>::nbits** = `13`, which is less than the specified register width (32). This is allowed and is equivalent to specifying a field of size 32 – 13 = 19 bits after `fld2[5]`. This reserved field cannot be accessed using **read()/write()** functions on the register object. In the numeric value passed to **write_val()** and in the return value of **read_val()**, the value of these bits is not defined by this standard.
5) If the name of an element in a register group is declared using an escaped identifier (see 4.3), the leading backslash is not considered as a part of the name. Therefore, the name string used in the above mentioned functions do not contain the leading backslash.

It is recommended to declare the register type as **pure**. This allows the PSS implementation to optimally handle large static register components.

21.14.2 PSS register group definition

A *register group* aggregates instances of registers and of other register groups.

The **reg_group_c** component is the base type for specifying register groups. Note that it is a **pure** component (see 9.1.7). It shall be illegal to extend the **reg_group_c** class.

Copyright © 2026 Accellera. All rights reserved.

```
struct node_s {
  string name;
  int    index;
};

pure component reg_group_c {
  pure function bit[64] get_offset_of_instance(string name);
  pure function bit[64] get_offset_of_instance_array(string name, int index);
  pure function bit[64] get_offset_of_path(list<node_s> path);
  solve function void set_handle(addr_handle_t addr);
  function addr_handle_t get_handle();
};
```

Syntax 163—PSS register group definition

A register group may instantiate registers and instances of other register groups. An instance of a register group may be created in another register group, or directly in a non-register-group component. In the latter case, the register group can be *associated* with an address region. The **`set_handle()`** function associates the register group with an address region. The definition of this function is implementation-defined. See 21.14.3 for more details on use of this function. The **`get_handle()`** function returns an address handle corresponding to the register group, independent of whether it is the root register group. See 21.14.4 for more details on obtaining an address handle to registers and register groups.

Each element in a register group (whether an instance of a register or an instance of another group) has a user-defined address offset relative to a notional base address of the register group.

The function **`get_offset_of_instance()`** retrieves the offset of a non-array element in a register group, by name of the element. The function **`get_offset_of_instance_array()`** retrieves the offset of an array element in a register group, by name of the element and index in the array.

For example, suppose `a` is an instance of a register group that has the following elements:

- A register instance, `r0`
- A register array instance, `r1[4]`

Calling `a.get_offset_of_instance("r0")` returns the offset of the element `r0`. Calling `a.get_offset_of_instance_array("r1", 2)` returns the offset at index `2` of element `r1`.

The function **`get_offset_of_path()`** retrieves the offset of a register from a hierarchical path of the register, starting from a given register group. The hierarchical path of the register is specified as a **list** of **`node_s`** objects. Each **`node_s`** object provides the name of the element (as a string) and an index (applicable if and only if the element is of array type). The first element of the list corresponds to an object directly instantiated in the given register group. Successive elements of the list correspond to an object instantiated in the register group referred by the predecessor node. The last element of the list corresponds to the final register instance.

For example, suppose `b` is an instance of a register group that has the following elements: a register group array instance `grp0[10],` which in turn has a register group instance `grp1`, which in turn has a register instance, `r0`. The hierarchical path of register `r0` in `grp1` within `grp0[5]` within `b` will then be the list (e.g., `path_to_r0`) with the following elements in succession:

Copyright © 2026 Accellera. All rights reserved.

- `[0]`: **`node_s`** object with **`name`** = `"grp0"` and **`index`** = 5
- `[1]`: **`node_s`** object with **`name`** = `"grp1"` (**`index`** is not used)
- `[2]`: **`node_s`** object with **`name`** = `"r0"` (**`index`** is not used)

Calling `b.get_offset_of_path(path_to_r0)` will return the offset of register `r0` relative to the base address of `b`.

For a given register group, users shall provide the implementation of either **`get_offset_of_path()`** or of both functions **`get_offset_of_instance()`** and **`get_offset_of_instance_array()`**. It shall be an error to provide an implementation of all three functions. These may be implemented as native PSS functions, or foreign-language binding may be used. These functions (when implemented) shall provide the relative offset of *all* the elements in the register group. These functions are called by a PSS tool to compute the offset for a register access (as described later in 21.14.4). Note that these functions are declared **pure** —the implementation shall not have side-effects.

Example 361 shows an example of a register group declaration.

```
pure component my_reg_grp0_c : reg_group_c {
   my_readonly_reg0_c   reg0;            // (1)
   my_reg1_c            reg1[4];         // (2)
   my_sub_reg_grp_c     sub;             // (3)
   reg_c<my_regx_s, WRITEONLY, 32> regx; // (4)

   // May be foreign, too
   function bit[64] get_offset_of_instance(string name) {
      match(name) {
        ["reg0"]:  return 0x0;
        ["sub"]:   return 0x20;
        ["regx"]:  return 0x0;             // (5)
        default: return -1; // Error case
      }
   }

   function bit[64] get_offset_of_instance_array(string name, int index) {
      match(name) {
        ["reg1"]:  return (0x4 + index*4);
        default: return -1; // Error case
      }
   }
}
```

Example 361—Example of register group declaration

Notes:

1) `my_readonly_reg0_c`, `my_reg1_c`, etc., are all register types (declarations not shown in the example).
2) Arrays of registers are allowed.
3) Groups may contain other groups (declaration of `my_sub_reg_grp_c` not shown in the example).
4) A direct instance of **`reg_c<>`** may be created in a register group.
5) Offsets of two elements may be same. A typical use case for this is when a **`READONLY`** and a **`WRITEONLY`** register share the same offset.

Copyright © 2026 Accellera. All rights reserved.

21.14.3 Association with address region

Before the read/write functions can be invoked on a register, the top-level register group (under which the register object has been instantiated) shall be associated with an address region, using the **`set_handle()`** function in that register group. This is done from within an **exec init_up** or **init_down** context. Only the top-level register group shall be associated with an address region; it shall be an error to call **`set_handle()`** on other register group instances. An example is shown in Example 362.

```
component my_component_c
{
   my_reg_grp0_c  grp0; // Top-level group

   transparent_addr_space_c<> sys_mem;

   exec init_up {
     transparent_addr_region_s<> mmio_region;
     addr_handle_t h;
     mmio_region.size = 1024;
     mmio_region.addr = 0xA0000000;

     h = sys_mem.add_nonallocatable_region(mmio_region);

     grp0.set_handle(h);
   }
};
```

Example 362—Top-level group and address region association

21.14.4 Obtaining register address handles

Registers are typically accessed via the **`read()`** and **`write()`** methods provided by the **`reg_c`** component (see 21.14.1). However, there are also cases where it is more convenient to access registers based on their addresses. The **`get_handle()`** method provided by the **`reg_c`** and **`reg_group_c`** components support these use cases.

Copyright © 2026 Accellera. All rights reserved.

```
component sub_blk : reg_group_c {
  reg_c<bit[32]>         r1;
  reg_c<bit[32]>         r2;
  reg_c<bit[32]>         r3;
  // ...
  reg_c<bit[32]>         r256;
}

component reg_blk : reg_group_c {
  sub_blk        s1;
  sub_blk        s2;
}

component pss_top {
  reg_blk       regs;

  action Entry {
    exec body {
        addr_handle_t s2_h = regs.s2.get_handle();
        repeat (ri : 256) {
            write32(make_handle_from_handle(s2_h, 4*ri), 0);
        }
    }
  }
}
```

Example 363—Accessing registers based on their addresses

The example above clears some large number of registers in a register block. These registers are enumerated for programmer convenience, but it is more convenient to clear them by address.

21.14.5 Translation of register read/write

The PSS implementation shall convert invocations of the register access functions described in Syntax 160 to invocations of the primitive read/write operations on the address associated with the register (see 21.13.9.1 and 21.13.9.2). The conversion shall proceed as follows:

a) The read/write function is selected based on the size of the register. For example, if the size of the register is 32, the function **`read32(addr_handle_t hndl)`** will be called for a register read.

b) The total offset is calculated by summing the offsets of all elements starting from the top-level register group to the register itself.

   1) If the function **`get_offset_of_path()`** is available in any intermediate register group instance, the PSS implementation will use that function to find the offset of the register relative to the register group.

   2) Otherwise, the function **`get_offset_of_instance_array()`** or **`get_offset_of_instance()`** is used, depending on whether or not the register instance or register group instance is an array.

   For example, in the expression (where `a`, `b`, `c`, and `d` are all instances of register groups and `reg` is a register object):

   ```
   comp.a.b.c.d[4].reg.write_val(10)
   ```

   if the function **`get_offset_of_path()`** is implemented in the type of element `c`, then the offset is calculated as:

Copyright © 2026 Accellera. All rights reserved.

```
offset = comp.a.get_offset_of_instance("b") +
         comp.a.b.get_offset_of_instance("c") +
         comp.a.b.c.get_offset_of_path(path)
```

where `path` is the **list** `[ {"d", 4}, {"reg", 0} ]`.

c) The handle for the access is calculated as **`make_handle_from_handle(h, offset)`**, where `h` is the handle set using **`set_handle()`** on the top-level register group.

21.14.6 Symbolic register names

Register instances (of type `reg_c<>`) within a register group hierarchy are currently accessed using address handles computed from the `get_offset_*` functions (see 21.14.2) and the handle set via `set_handle()` (see 21.14.3). The resulting generated target code refers to registers by their resolved numeric address values.

It is sometimes beneficial to refer to a register by a symbolic name rather than a numeric address. This can facilitate compile-time type checking, improve readability of generated code for software developers, and allow a PSS model to align with symbolic register definitions already present in an existing software stack.

This section defines extensions to `reg_group_c` and to `addr_reg_pkg` that enable symbolic register access as an alternative to address-handle-based access.

21.14.6.1 Mnemonic functions on reg_group_c

Three new functions are added to `reg_group_c` alongside the existing `get_offset_*` and `set_handle()` declarations, together with a new `set_mnemonic()` function as shown in Syntax 164.

```
pure component reg_group_c {
  // Existing declarations (21.14.2, 21.14.3):
  pure function bit[64] get_offset_of_instance(string name);
  pure function bit[64] get_offset_of_instance_array(string name, int index);
  pure function bit[64] get_offset_of_path(list<node_s> path);
  solve function void   set_handle(addr_handle_t addr);

  // New additions:
  solve pure function string get_mnemonic_of_instance(string name);
  solve pure function string get_mnemonic_of_instance_array(string name,
                                                            int index);
  solve pure function string get_mnemonic_of_path(list<node_s> path);
  solve function void       set_mnemonic(string prefix);
};
```

Syntax 164—Mnemonic functions on reg_group_c

— `get_mnemonic_of_instance(name)`: Returns the mnemonic fragment to be appended when traversing to a non-array child element with the specified name.

— `get_mnemonic_of_instance_array(name, index)`: Returns the mnemonic fragment to be appended when traversing to an element index of the child array with the specified name.

Copyright © 2026 Accellera. All rights reserved.

— `get_mnemonic_of_path(path)`: Returns the mnemonic fragment for the sub-hierarchy rooted at this group and described by the path, a list of `node_s` objects. The semantics of path match those of the corresponding `get_offset_of_path()` argument (see 21.14.2).

— `set_mnemonic(prefix)`: Sets a root-level prefix string for the top-level register group. When non-empty, this prefix is prepended to the concatenated mnemonic before the first

— `get_mnemonic_*` contribution. It may only be called on a top-level register group instance (i.e., on the same instance that `set_handle()` is called on).

Rules:

a) For a given register group, users shall provide the implementation of either `get_mnemonic_of_path()` or of both `get_mnemonic_of_instance()` and `get_mnemonic_of_instance_array()`. It shall be an error to provide an implementation of all three functions.

b) The argument semantics of each mnemonic function are identical to those of the corresponding `get_offset_*` function in the same component type (see 21.14.2).

c) These functions are declared **solve pure** and shall not have side-effects.

d) `set_mnemonic()` shall only be called in **exec init_down** or **exec init_up** blocks. It shall be an error to call `set_mnemonic()` on a register group that is not a top-level group (i.e., on a group that is a child of another register group).

21.14.6.2 Syntax

```
package addr_reg_pkg {
  solve function void use_symbolic_reg_names(ref reg_group_c grp, bool enable);
}
```

Syntax 165—Use_symbolic_reg_names function

`use_symbolic_reg_names()` configures the register access mode for the top-level register group instance *grp* (see Syntax 165). It shall only be called in **exec init_down** or **exec init_up** blocks. *grp* shall be a top-level register group instance; it shall be an error to call this function on a non-top-level group instance.

The access mode controls how the PSS tool generates code for register read/write invocations across the entire component hierarchy under *grp*:

a) When called with `enable == true`, the PSS tool shall use symbolic names for all register access code generated under *grp*. It shall be an error if any register group in the hierarchy under *grp* does not provide an implementation of the `get_mnemonic_*` functions.

b) When called with `enable == false`, the PSS tool shall use address handles (derived from `get_offset_*` functions and the handle set via `set_handle()`) for all register access code generated under *grp*. It shall be an error if any register group in the hierarchy under *grp* does not provide an implementation of the `get_offset_*` functions.

c) If `use_symbolic_reg_names()` is not called for a given top-level register group instance, address-handle mode is used by default.

NOTE—PSS tool implementations may additionally support a mode in which both `get_offset_*` and `get_mnemonic_*` functions are invoked for the same register access, with the resolved address used for the actual access and the symbolic name emitted as an annotation (e.g., a comment in generated C code). This is tool-specific behavior and is not mandated by this specification.

Copyright © 2026 Accellera. All rights reserved.

21.14.6.3 Translation of register access using symbolic names

When symbolic mode is enabled for a top-level register group, the PSS tool constructs a unique symbolic name for each register access and uses it in place of the address handle argument to the underlying primitive read/write operations.

The symbolic name is constructed as follows:

a) Start with the prefix string set by `set_mnemonic()` on the top-level group. If `set_mnemonic()` was not called, the prefix is the empty string.

b) Traverse the component hierarchy from the top-level register group down to the target register. At each level:

   1) If `get_mnemonic_of_path()` is implemented in the current group, invoke it with the `node_s` path representing the remainder of the hierarchy below that group. Append the returned string and stop descending.

   2) Otherwise, invoke `get_mnemonic_of_instance()` or `get_mnemonic_of_instance_array()` as appropriate for the child element being traversed. Append the returned string and continue to the next level.

c) The final symbolic name is the concatenation of the prefix and all appended fragments, in top-down order. No additional delimiter characters are inserted between fragments; the user implementation is fully responsible for the content of each fragment, including any required separator characters.

d) The PSS tool translates a register access expression such as `comp.grp.sub.reg.write_val(v)` to the equivalent call using the constructed symbolic name directly in place of the address handle argument. The specific form of the generated code is tool-dependent.

21.14.6.4 Examples

21.14.6.4.1 Example A: Single register group with flat symbolic names

Suppose the target platform (C-based) defines symbolic register names for a DMA block (see Example 364).

```
const int* DMA_REG_A;
const int* DMA_REG_B;
const int* DMA_REG_C[3];
```

Example 364—Target platform symbolic register names

The corresponding PSS register group implements the mnemonic functions to return these names (see Example 365).

Copyright © 2026 Accellera. All rights reserved.

```
pure component dma_reg_group_c : reg_group_c {
    reg_c<R1_s, READWRITE, 32> reg_a;
    reg_c<R1_s, READWRITE, 32> reg_b;
    reg_c<R1_s, READWRITE, 32> reg_c[3];

    solve pure function string get_mnemonic_of_instance(string name) {
        if (name == "reg_a") { return "DMA_REG_A"; }
        if (name == "reg_b") { return "DMA_REG_B"; }
    }

    solve pure function string get_mnemonic_of_instance_array(string name,
                                                              int index) {
        if (name == "reg_c") { return format("DMA_REG_C[%d]", index); }
    }
}
```

Example 365—Register group with flat symbolic names

For a write access `comp.regs.reg_c[2].write_val(v)`, the PSS tool constructs the symbolic name `"DMA_REG_C[2]"` and generates the corresponding target call.

21.14.6.4.2 Example B: Hierarchical symbolic names with explicit delimiter

Assume two instances of the DMA module above (see Example 364). The target platform (C) defines symbolic names using a struct containing each instance's registers (see Example 366).

```
typedef struct dma_regs {
    int * DMA_REG_A;
    int * DMA_REG_B;
    int * DMA_REG_C[3];
} dma_regs;
static const dma_regs DMA_0 = { ... };
static const dma_regs DMA_1 = { ... };
```

Example 366—Target platform struct-based symbolic names

To access a register, the symbolic name takes the form `DMA_0.DMA_REG_B`. A containing register group produces the struct-instance prefix; the mnemonic functions in `dma_reg_group_c` from Example A (see 21.14.6.4.1) produce the field name (see Example 367).

Copyright © 2026 Accellera. All rights reserved.

```
pure component dma_cluster_group_c : reg_group_c {
    dma_reg_group_c dma0, dma1;

    solve pure function string get_mnemonic_of_instance(string name) {
        if (name == "dma0") { return "DMA_0."; }
        // delimiter included in fragment
        if (name == "dma1") { return "DMA_1."; }
    }

    solve pure function string get_mnemonic_of_instance_array(string name,
                                                              int index) {
        // No arrays in this group
    }
}
```

Example 367—Hierarchical register group with struct-accessor delimiter

For a write access to `dma1.reg_b`, the PSS tool concatenates `"DMA_1."` (from `dma_cluster_group_c`) with `"DMA_REG_B"` (from `dma_reg_group_c`) to produce `"DMA_1.DMA_REG_B"`.

Note that the `'.'` delimiter is explicitly embedded in the fragment string returned by `get_mnemonic_of_instance()` in `dma_cluster_group_c`. Because the PSS tool performs concatenation without inserting any separators, the user has full control over the resulting format. Alternative formats such as `DMA_1->DMA_REG_A` or `DMA_1+DMA_REG_A` are achieved by returning `"DMA_1->"` or `"DMA_1+"`, respectively, from the mnemonic function.

21.14.6.4.3 Example C: Recursive hierarchy with enable mechanism

Example 368 shows a two-level register hierarchy with two DMA engine instances.

Copyright © 2026 Accellera. All rights reserved.

```
static const int NUM_REG_C = 3;
static const int NUM_ENG   = 2;

struct R1_s : packed_s<> { bit[32] fld; }

pure component dma_reg_group_c : reg_group_c {
    reg_c<R1_s, READWRITE, 32> reg_a;
    reg_c<R1_s, READWRITE, 32> reg_b;
    reg_c<R1_s, READWRITE, 32> reg_c[NUM_REG_C];

    solve pure function string get_mnemonic_of_instance(string name) {
        if (name == "reg_a") { return "DMA_REG_A"; }
        if (name == "reg_b") { return "DMA_REG_B"; }
    }

    solve pure function string get_mnemonic_of_instance_array(string name,
                                                              int index) {
        if (name == "reg_c") { return format("DMA_REG_C[%d]",
                                             index % NUM_REG_C); }
    }
}
pure component subsys_reg_group_c : reg_group_c {
    dma_reg_group_c eng[NUM_ENG];

    solve pure function string get_mnemonic_of_instance(string name) {
        // No non-array children in this group
    }

    solve pure function string get_mnemonic_of_instance_array(string name,
                                                              int index) {
        if (name == "eng") {
            if ((index % 2) == 0) { return "DMA0_"; }
            if ((index % 2) == 1) { return "DMA1_"; }
        }
    }
}

extend component pss_top {
    subsys_reg_group_c         subsys_regs;
    transparent_addr_space_c<> mem;

    exec init_down {
        transparent_addr_region_s<> mmio;
        addr_handle_t h;
        mmio.size = 0x80000;
        mmio.addr = 0xA0000000;
        h = mem.add_nonallocatable_region(mmio);
        subsys_regs.set_handle(h);
        subsys_regs.set_mnemonic("");           // no root prefix
        use_symbolic_reg_names(subsys_regs, true);
    }
```

Example 368—Full hierarchical example with symbolic name enable

Copyright © 2026 Accellera. All rights reserved.

```
        action test_basic {
            exec body {
                comp.subsys_regs.eng[0].reg_a.write_val(0);    // DMA0_DMA_REG_A
                comp.subsys_regs.eng[1].reg_c[1].write_val(0);
                // DMA1_DMA_REG_C[1]
            }
        }
    }
```

Example 368—Full hierarchical example with symbolic name enable (Cont.)

The PSS tool concatenates the mnemonic fragments as shown in Table 34.

Table 34—Mnemonic construction for hierarchical register access

Access expression	Fragment from subsys_reg_group_c	Fragment from dma_reg_group_c	Resulting symbolic name
`eng[0].reg_a`	`"DMA0_"`	`"DMA_REG_A"`	`"DMA0_DMA_REG_A"`
`eng[0].reg_c[2]`	`"DMA0_"`	`"DMA_REG_C[2]"`	`"DMA0_DMA_REG_C[2]"`
`eng[1].reg_b`	`"DMA1_"`	`"DMA_REG_B"`	`"DMA1_DMA_REG_B"`
`eng[1].reg_c[1]`	`"DMA1_"`	`"DMA_REG_C[1]"`	`"DMA1_DMA_REG_C[1]"`

21.14.6.5 Interaction with get_offset_* functions

The `get_mnemonic_*` and `get_offset_*` function sets are independent. A register group may implement both sets simultaneously. The active set for a given top-level group instance is determined by the `use_symbolic_reg_names()` configuration (see Table 35).

Table 35—Register access mode selection

Use_symbolic_reg_names() called?	Enable value	Mode in effect	Required implementation
No (default)	–	Address-handle mode	`get_offset_*` functions
Yes	`false`	Address-handle mode	`get_offset_*` functions
Yes	`true`	Symbolic mode	`get_mnemonic_*` functions

It shall be an error for a register group to implement both `get_mnemonic_of_path()` and either of `get_mnemonic_of_instance()` or `get_mnemonic_of_instance_array()`. This rule applies regardless of which access mode is currently active.

Copyright © 2026 Accellera. All rights reserved.

21.14.7 Recommended packaging

It is recommended that all the register (and register group) definitions of a device be placed in a separate file and in a separate package by themselves, as shown in Example 369.

```
// In my_IP_regs.pss
package my_IP_regs {
   import addr_reg_pkg::*;
   struct my_reg0_s : packed_s<> { ... };
   pure component my_reg0_c : reg_c<my_reg0_s, READWRITE, 32> { ... };
   // ... etc: other registers

   pure component my_reg_group_c : reg_group_c {
      my_reg0_c r0;
      // ... etc: other registers
   };
}
```

Example 369—Recommended packaging

This ensures that the register file can be easily generated from a register specification (e.g., IP-XACT).

Copyright © 2026 Accellera. All rights reserved.

Annex A
(informative)
Bibliography

[B1] IEEE 100, *The Authoritative Dictionary of IEEE Standards Terms,* Seventh Edition. New York: Institute of Electrical and Electronics Engineers, Inc.

Copyright © 2026 Accellera. All rights reserved.

Annex B
(normative)
Formal syntax

The PSS formal syntax is described using Backus-Naur Form (BNF). The syntax of the PSS source is derived from the starting symbol `Model`. If there is a conflict between a grammar element shown anywhere in this standard and the material in this annex, the material shown in this annex shall take precedence.

```
Model ::= { portable_stimulus_description }

portable_stimulus_description ::=
      package_body_item
    | package_declaration
    | component_declaration
```

B.1 Package declarations

```
package_declaration ::= package package_id_path  { { package_body_item } }

package_id_path ::= package_identifier { :: package_identifier }

package_body_item ::=
      abstract_action_declaration
    | abstract_monitor_declaration
    | struct_declaration
    | enum_declaration
    | covergroup_declaration
    | generic_constraint_declaration
    | function_decl
    | import_class_decl
    | procedural_function
    | import_function
    | target_template_function
    | export_action
    | typedef_declaration
    | import_stmt
    | extend_stmt
    | const_field_declaration
    | component_declaration
    | package_declaration
    | compile_assert_stmt
    | package_body_compile_if
    | annotation_declaration
    | annotation_usage package_body_item
    | export_function
    | stmt_terminator

import_stmt ::= import package_import_pattern ;

package_import_pattern ::= type_identifier [ package_import_qualifier ]

package_import_qualifier ::=
      package_import_wildcard
    | package_import_alias
```

Copyright © 2026 Accellera. All rights reserved.

```
package_import_wildcard ::= :: *

package_import_alias ::= as package_identifier

extend_stmt ::=
      extend action type_identifier { { action_body_item } }
    | extend component type_identifier { { component_body_item } }
    | extend struct_kind type_identifier { { struct_body_item } }
    | extend enum type_identifier { [ enum_item { , enum_item } ] }
    | extend annotation type_identifier { { annotation_body_item } }

const_field_declaration ::= [ static ] const data_declaration

stmt_terminator ::= ;

annotation_declaration ::= annotation annotation_identifier
    [ template_param_decl_list ]
    [ annotation_super_spec ] { { annotation_body_item } }

annotation_super_spec ::= : type_identifier

annotation_body_item ::=
  annotation_attr_field
  | compile_assert_stmt
  | annotation_body_compile_if
  | annotation_usage annotation_body_item
  | stmt_terminator

annotation_attr_field ::= [static const] data_declaration

annotation_body_compile_if ::=
  compile if ( constant_expression ) annotation_body_compile_if_item
    [ else annotation_body_compile_if_item ]

annotation_body_compile_if_item ::= { { annotation_body_item } }

annotation_usage ::= @ annotation_type_identifier [ annotation_params_list ]
annotation_params_list ::= { annotation_param_item {, annotation_param_item} }
annotation_param_item ::= . identifier = constant_expression
```

B.2 Action declarations

```
action_declaration ::= action action_identifier [ template_param_decl_list ]
    [ action_super_spec ] { { action_body_item } }

abstract_action_declaration ::= abstract action_declaration

override_action_declaration ::= override action action_identifier
    { { action_body_item } }

action_super_spec ::= : type_identifier

action_body_item ::=
      activity_declaration
    | override_declaration
```

Copyright © 2026 Accellera. All rights reserved.

```
    | constraint_declaration
    | action_field_declaration
    | symbol_declaration
    | covergroup_declaration
    | generic_constraint_declaration
    | exec_block_stmt
    | activity_scheduling_constraint
    | attr_group
    | compile_assert_stmt
    | covergroup_instantiation
    | action_body_compile_if
    | annotation_usage action_body_item
    | stmt_terminator

activity_declaration ::= activity { { activity_stmt } }

action_field_declaration ::=
      attr_field
    | activity_data_field
    | action_handle_declaration
    | object_ref_field_declaration

object_ref_field_declaration ::=
      flow_ref_field_declaration
    | resource_ref_field_declaration

flow_ref_field_declaration ::=
    ( input | output ) flow_object_type object_ref_field { , object_ref_field } ;

resource_ref_field_declaration ::=
   ( lock | share ) resource_object_type object_ref_field {, object_ref_field } ;

flow_object_type ::=
      buffer_type_identifier
    | state_type_identifier
    | stream_type_identifier

resource_object_type ::= resource_type_identifier

object_ref_field ::= identifier { array_dim }

action_handle_declaration ::= action_type_identifier action_instantiation ;

action_instantiation ::= action_handle_instance {,
        action_handle_instance} ;

action_handle_instance ::=
    action_handle_array_instance
    | action_handle_single_instance

action_handle_array_instance ::=
    action_handle_identifier { array_dim }
    action_handle_single_instance ::=
    action_handle_identifier [ action_initializer_list ]

action_initializer_list ::= { action_initializer {, action_initializer} }

action_initializer ::= . hierarchical_id = expression
```

Copyright © 2026 Accellera. All rights reserved.

```
activity_data_field ::= action data_declaration

activity_scheduling_constraint ::= constraint ( parallel | sequence )
      { hierarchical_id , hierarchical_id { , hierarchical_id } };
```

B.3 Struct declarations

```
struct_declaration ::= struct_kind struct_identifier
    [ template_param_decl_list ] [ struct_super_spec ] { { struct_body_item } }

struct_kind ::=
      struct
    | object_kind

object_kind ::=
      buffer
    | stream
    | state
    | resource

struct_super_spec ::= : type_identifier

struct_body_item ::=
      constraint_declaration
    | attr_field
    | typedef_declaration
    | generic_constraint_declaration
    | exec_block_stmt
    | attr_group
    | compile_assert_stmt
    | covergroup_declaration
    | covergroup_instantiation
    | struct_body_compile_if
    | annotation_usage struct_body_item
    | stmt_terminator
```

B.4 Exec blocks

```
exec_block_stmt ::=
  exec_block
  | target_code_exec_block
  | target_file_exec_block

exec_block ::= exec exec_kind { { procedural_stmt } }

exec_kind ::=
  pre_solve
  | post_solve
  | pre_body
  | body
  | header
  | declaration
  | run_start
```

Copyright © 2026 Accellera. All rights reserved.

```
  | run_end
  | init_down
  | init_up

exec_super_stmt ::= super ;

target_code_exec_block ::= exec exec_kind language_identifier =
                               [exec_block_tag :] string_literal ;

target_file_exec_block ::= exec file filename_string =
                               [exec_block_tag :] string_literal ;

exec_block_tag ::= type_identifier [struct_literal]
```

B.5 Functions

```
procedural_function ::= [ platform_qualifier ] [ pure ] [ static ] function
   function_prototype { { procedural_stmt } }

function_decl ::= [ platform_qualifier ][ pure ] [ static ] function
   function_prototype ;
platform_qualifier ::=
     target
   | solve

function_prototype ::=
   function_return_type function_identifier function_parameter_list_prototype

function_return_type ::=
      void
    | data_type

function_parameter_list_prototype ::=
     ( [ function_parameter { , function_parameter } ] )
   | ( { function_parameter , } varargs_parameter )

function_parameter ::=
    [ function_parameter_dir | const ]
      data_type identifier [ = constant_expression ]
   | [const] ( type | ref ref_type_category | plain_type_category  ) identifier

function_parameter_dir ::=
      input
    | output
    | inout

varargs_parameter ::=
      ( data_type | type | ref ref_type_category | plain_type_category ) ...
   identifier
```

B.6 Foreign procedural interface

```
import_function ::=
       import [ platform_qualifier ] [ language_identifier ]
```

Copyright © 2026 Accellera. All rights reserved.

```
        function type_identifier ;
    | import [ platform_qualifier ] [ language_identifier ] [ static ]
        function function_prototype ;

platform_qualifier ::=
      target
    | solve

target_template_function ::=
      target language_identifier [ static ]
        function function_prototype = string_literal ;

import_class_decl ::= import class import_class_identifier
   [ import_class_extends ] { { import_class_function_decl } }

import_class_extends ::= : type_identifier { , type_identifier }

import_class_function_decl ::= function_prototype ;

export_action ::= export [ platform_qualifier ] action_type_identifier
   function_parameter_list_prototype ;

export_function ::= export target function type_identifier ;
```

B.7 Procedural statements

```
procedural_stmt ::=
  procedural_sequence_block_stmt
  | procedural_data_declaration
  | procedural_assignment_stmt
  | procedural_void_function_call_stmt
  | procedural_return_stmt
  | procedural_repeat_stmt
  | procedural_foreach_stmt
  | procedural_if_else_stmt
  | procedural_match_stmt
  | procedural_break_stmt
  | procedural_continue_stmt
  | procedural_randomization_stmt
  | procedural_compile_if
  | procedural_yield_stmt
  | exec_super_stmt
  | annotation_usage procedural_stmt
  | stmt_terminator

procedural_sequence_block_stmt ::= [ sequence ] { { procedural_stmt } }

procedural_data_declaration ::= data_type procedural_data_instantiation
  { , procedural_data_instantiation } ;

procedural_data_instantiation ::= identifier { array_dim } [ = expression ]

procedural_assignment_stmt ::= ref_path assign_op expression ;

procedural_void_function_call_stmt ::= [ ( void ) ] function_call ;
```

Copyright © 2026 Accellera. All rights reserved.

```
procedural_return_stmt ::= return [ expression ] ;

procedural_repeat_stmt ::=
  repeat ( [ index_identifier : ] expression ) procedural_stmt
  | repeat procedural_stmt while ( expression ) ;
  | while ( expression ) procedural_stmt

procedural_foreach_stmt ::=
  foreach ( [ iterator_identifier : ] expression  [ [ index_identifier ] ] )
    procedural_stmt

procedural_if_else_stmt ::=
  if ( expression ) procedural_stmt [ else procedural_stmt ]

procedural_match_stmt ::=
  match ( match_expression )
    { procedural_match_choice { procedural_match_choice } }

procedural_match_choice ::=
    [ open_range_list ] : procedural_stmt
  | default : procedural_stmt

procedural_break_stmt ::= break ;

procedural_continue_stmt ::= continue ;

procedural_randomization_stmt ::=
  randomize procedural_randomization_target procedural_randomization_term

procedural_randomization_target ::= hierarchical_id { , hierarchical_id }

procedural_randomization_term ::=
    with constraint_set
  | ;
procedural_yield_stmt ::=
  yield ;
```

B.8 Component declarations

```
component_declaration ::=
    [ pure ] component component_identifier [ template_param_decl_list ]
      [ component_super_spec ] { { component_body_item } }

component_super_spec ::= : type_identifier

component_body_item ::=
      override_declaration
    | component_data_declaration
    | component_pool_declaration
    | action_declaration
    | abstract_action_declaration
    | override_action_declaration
    | generic_constraint_declaration
    | object_bind_stmt
```

Copyright © 2026 Accellera. All rights reserved.

```
    | exec_block
    | struct_declaration
    | enum_declaration
    | covergroup_declaration
    | function_decl
    | import_class_decl
    | procedural_function
    | import_function
    | export_function
    | target_template_function
    | export_action
    | typedef_declaration
    | import_stmt
    | extend_stmt
    | compile_assert_stmt
    | attr_group
    | component_body_compile_if
    | monitor_declaration
    | cover_stmt
    | annotation_usage component_body_item
    | stmt_terminator

component_data_declaration ::=
    [ access_modifier ] [ component_data_decl_qualifier ] data_declaration
component_data_decl_qualifier ::=
  static const
 | mutable
 | instance

component_pool_declaration ::=
    pool [ [ expression ] ] type_identifier identifier ;

object_bind_stmt ::= bind hierarchical_id object_bind_item_or_list ;

object_bind_item_or_list ::=
      object_bind_item_path
    | { object_bind_item_path { , object_bind_item_path } }

object_bind_item_path ::= { component_path_elem . } object_bind_item

component_path_elem ::= component_identifier { [ domain_open_range_list ] }

object_bind_item ::=
      action_type_identifier . identifier { [ domain_open_range_list ] }
    | *
```

B.9 Activity statements

```
activity_stmt ::=
      [ label_identifier : ] labeled_activity_stmt
    | activity_action_traversal_stmt
    | activity_data_field
    | activity_bind_stmt
    | action_handle_declaration
    | activity_constraint_stmt
    | activity_super_stmt
```

Copyright © 2026 Accellera. All rights reserved.

```
    | activity_scheduling_constraint
    | annotation_usage activity_stmt
    | stmt_terminator

labeled_activity_stmt ::=
      activity_sequence_block_stmt
    | activity_parallel_stmt
    | activity_schedule_stmt
    | activity_repeat_stmt
    | activity_foreach_stmt
    | activity_select_stmt
    | activity_if_else_stmt
    | activity_match_stmt
    | activity_replicate_stmt
    | activity_atomic_block_stmt
    | symbol_call

activity_action_traversal_stmt ::=
    action_handle_traversal_stmt
    | action_type_traversal_stmt

action_handle_traversal_stmt ::=
      identifier { [ expression ] } [ action_initializer_list ]
          inline_constraints_or_empty

action_type_traversal_stmt ::=
      [ label_identifier : ] do type_identifier [ action_initializer_list ]
          inline_constraints_or_empty

inline_constraints_or_empty ::=
      action_activity_inline_constraints_or_empty
    | monitor_activity_inline_constraints_or_empty
action_activity_inline_constraints_or_empty ::=
     with constraint_set
    | ;

activity_sequence_block_stmt ::= [ sequence ] { { activity_stmt } }

activity_parallel_stmt ::= parallel [ activity_join_spec ] { { activity_stmt } }

activity_schedule_stmt ::= schedule [ activity_join_spec ] { { activity_stmt } }

activity_join_spec ::=
      activity_join_branch
    | activity_join_select
    | activity_join_none
    | activity_join_first

activity_join_branch ::= join_branch ( label_identifier { , label_identifier } )

activity_join_select ::= join_select ( expression )

activity_join_none ::= join_none

activity_join_first ::= join_first ( expression )

activity_repeat_stmt ::=
      repeat ( [ index_identifier : ] expression ) activity_stmt
```

Copyright © 2026 Accellera. All rights reserved.

```
    | repeat activity_stmt while ( expression ) ;

activity_foreach_stmt ::= foreach ( [ iterator_identifier : ] expression
    [ [ index_identifier ] ] ) activity_stmt

activity_select_stmt ::= select { select_branch select_branch { select_branch } }

select_branch ::= [[ ( expression ) ][ [ expression ] ] : ] activity_stmt

activity_if_else_stmt ::= if ( expression ) activity_stmt [ else activity_stmt ]

activity_match_stmt ::=
    match ( match_expression ) { match_choice { match_choice } }

match_expression ::= expression

match_choice ::=
      [ open_range_list ] : activity_stmt
    | default : activity_stmt

activity_replicate_stmt ::= replicate ( [ index_identifier : ] expression )
    [ label_identifier [ ] : ] labeled_activity_stmt

activity_super_stmt ::= super ;

activity_atomic_block_stmt ::= atomic { { activity_stmt } }

activity_bind_stmt ::= bind hierarchical_id activity_bind_item_or_list ;

activity_bind_item_or_list ::=
      hierarchical_id
    | { hierarchical_id_list }

activity_constraint_stmt ::= constraint constraint_set

symbol_declaration ::=
    symbol symbol_identifier [ ( symbol_paramlist ) ] { { activity_stmt } }

symbol_paramlist ::= [ symbol_param { , symbol_param } ]

symbol_param ::= data_type identifier
```

B.10 Overrides

```
override_declaration ::= override { { override_stmt } }

override_stmt ::=
      type_override
    | instance_override
    | override_compile_if
    | annotation_usage override_stmt
    | stmt_terminator

type_override ::= type type_identifier with type_identifier ;
```

Copyright © 2026 Accellera. All rights reserved.

```
instance_override ::= instance hierarchical_id with type_identifier ;
```

B.11 Data declarations

```
data_declaration ::= data_type data_instantiation { , data_instantiation } ;

data_instantiation ::= identifier { array_dim } [ = constant_expression ]

array_dim ::= [ constant_expression ]

attr_field ::= [ access_modifier ] [ rand | static const ] data_declaration

access_modifier ::= public | protected | private

attr_group ::= access_modifier :
```

B.12 Template types

```
template_param_decl_list ::= < template_param_decl { , template_param_decl } >

template_param_decl ::= type_param_decl | value_param_decl

type_param_decl ::= generic_type_param_decl | category_type_param_decl

generic_type_param_decl ::= type identifier [ = type_identifier ]

category_type_param_decl ::=
    type_category identifier [ type_restriction ] [ = type_identifier ]

type_restriction ::= : type_identifier

value_param_decl ::= data_type identifier [ = constant_expression ]

template_param_value_list ::=
    < [ template_param_value { , template_param_value } ] >

template_param_value ::= constant_expression | data_type
```

B.13 Data types

```
data_type ::=
      scalar_data_type
    | collection_type
    | reference_type
    | type_identifier

scalar_data_type ::=
      chandle_type
    | integer_type
    | string_type
    | bool_type
    | enum_type
    | float_type
```

Copyright © 2026 Accellera. All rights reserved.

```
casting_type ::=
      integer_type
    | bool_type
    | enum_type
    | float_type
    | reference_type
    | type_identifier

chandle_type ::= chandle

integer_type ::= integer_atom_type
    [ [ constant_expression ] ]
    [ in [ domain_open_range_list ] ]

integer_atom_type ::=
      int
    | bit

domain_open_range_list ::=
    domain_open_range_value { , domain_open_range_value }

domain_open_range_value ::=
      constant_expression [ .. constant_expression ]
    | constant_expression ..
    | .. constant_expression

string_type ::= string [ in [ string_literal { , string_literal } ] ]

bool_type ::= bool

enum_declaration ::=
     enum enum_identifier [ : data_type ] { [ enum_item { , enum_item } ] }

enum_item ::= identifier [ = constant_expression ]

enum_type ::= enum_type_identifier [ in [ domain_open_range_list ] ]

float_type ::=
      float32
    | float64

collection_type ::=
      array < data_type , array_size_expression >
    | list < data_type >
    | map < data_type , data_type >
    | set < data_type >

array_size_expression ::= constant_expression

reference_type ::= ref entity_type_identifier

typedef_declaration ::= typedef data_type identifier ;
ref_type_category ::=
      action
    | monitor
```

Copyright © 2026 Accellera. All rights reserved.

```
      | component
      | object_kind

plain_type_category ::=
        struct
      | numeric

type_category ::=
        ref_type_category
      | plain_type_category
```

B.14 Constraints

```
constraint_declaration ::=
        constraint constraint_set
      | [ dynamic ] constraint identifier constraint_block
      | generic_constraint_declaration

constraint_set ::=
        constraint_body_item
      | constraint_block

constraint_block ::= { { constraint_body_item } }

constraint_body_item ::=
        expression_constraint_item
      | foreach_constraint_item
      | forall_constraint_item
      | if_constraint_item
      | implication_constraint_item
      | unique_constraint_item
      | default hierarchical_id == constant_expression ;
      | default disable hierarchical_id ;
      | dist_directive
      | constraint_body_compile_if
      | soft_constraint_item
      | annotation_usage constraint_body_item3
      | stmt_terminator

expression_constraint_item ::= expression ;

foreach_constraint_item ::=
     foreach ( [ iterator_identifier : ] expression [ [ index_identifier ] ] )
        constraint_set

forall_constraint_item ::=
    forall ( iterator_identifier : type_identifier [ in ref_path ] ) constraint_set

if_constraint_item ::= if ( expression ) constraint_set [ else constraint_set ]

implication_constraint_item ::= expression -> constraint_set
soft_constraint_item ::= soft expression ;
unique_constraint_item ::= unique unique_constraint_argument ;

unique_constraint_argument ::=
        hierarchical_id [ [ array_slice ] ]
```

Copyright © 2026 Accellera. All rights reserved.

```
      | { hierarchical_id_list } ;

dist_directive ::= dist expression in [ dist_list ] ;

dist_list ::= dist_item { , dist_item }

dist_item ::= open_range_value [ dist_weight ]

dist_weight ::=
     := expression
   | :/ expression

generic_constraint_declaration ::=
      generic_constraint_bool
    | generic_constraint_value

generic_constraint_bool ::=
    [static] constraint identifier generic_constraint_params constraint_set

generic_constraint_value ::=
    [static] constraint generic_constraint_data_type identifier
    generic_constraint_params expression_constraint_item ;

generic_constraint_params ::=
    ([generic_constraint_param {, generic_constraint_param}])

generic_constraint_param ::=
    [const] generic_constraint_data_type identifier

generic_constraint_data_type ::=
      numeric
    | data_type
```

B.15 Data coverage specification

```
covergroup_declaration ::= covergroup covergroup_identifier
   ( covergroup_port { , covergroup_port } ) { { covergroup_body_item } }

covergroup_port ::= data_type identifier

covergroup_body_item ::=
      covergroup_option
    | covergroup_coverpoint
    | covergroup_cross
    | covergroup_body_compile_if
    | annotation_usage covergroup_body_item
    | stmt_terminator

covergroup_option ::=
      option . identifier = constant_expression ;

covergroup_instantiation ::=
      covergroup_type_instantiation
    | inline_covergroup

inline_covergroup ::= covergroup { { covergroup_body_item } } identifier ;
```

Copyright © 2026 Accellera. All rights reserved.

```
covergroup_type_instantiation ::=
    covergroup_type_identifier covergroup_identifier
      ( covergroup_portmap_list ) covergroup_options_or_empty

covergroup_portmap_list ::=
       covergroup_portmap { , covergroup_portmap }
     | hierarchical_id_list

covergroup_portmap ::= . identifier ( hierarchical_id )

covergroup_options_or_empty ::=
       with { { covergroup_option } }
     | ;
covergroup_coverpoint ::= [ [ data_type ] coverpoint_identifier : ] coverpoint
    expression [ iff ( expression ) ] bins_or_empty

bins_or_empty ::=
       { { covergroup_coverpoint_body_item } }
     | ;

covergroup_coverpoint_body_item ::=
       covergroup_option
     | covergroup_coverpoint_binspec

covergroup_coverpoint_binspec ::= bins_keyword identifier
    [ [ [ constant_expression ] ] ] = coverpoint_bins

coverpoint_bins ::=
       [ covergroup_range_list ] [ with ( covergroup_expression ) ] ;
     | coverpoint_identifier with ( covergroup_expression ) ;
     | default ;

covergroup_range_list ::= covergroup_value_range { , covergroup_value_range }

covergroup_value_range ::=
       expression
     | expression .. [ expression ]
     | [ expression ] .. expression

bins_keyword ::= bins | illegal_bins | ignore_bins

covergroup_expression ::= expression

covergroup_cross ::=
     covercross_identifier : cross coverpoint_identifier
       { , coverpoint_identifier }[ iff ( expression ) ] cross_item_or_null

cross_item_or_null ::=
       { { covergroup_cross_body_item } }
     | ;

covergroup_cross_body_item ::=
       covergroup_option
     | covergroup_cross_binspec
```

Copyright © 2026 Accellera. All rights reserved.

```
covergroup_cross_binspec ::= bins_keyword identifier = covercross_identifier
    with ( covergroup_expression ) ;
```

B.16 Behavioral coverage specification

```
cover_stmt::=
  [ label_identifier: ] cover type_identifier ;
| [ label_identifier: ] cover { { monitor_body_item } }

monitor_declaration ::= monitor monitor_identifier
  [ template_param_decl_list ] [ monitor_super_spec ] { { monitor_body_item  } }

abstract_monitor_declaration ::= abstract monitor_declaration

monitor_super_spec ::= : type_identifier

monitor_body_item ::=
    monitor_activity_declaration
  | override_declaration
  | monitor_constraint_declaration
  | monitor_field_declaration
  | covergroup_declaration
  | attr_group
  | compile_assert_stmt
  | covergroup_instantiation
  | monitor_body_compile_if
  | annotation_usage monitor_body_item
  | stmt_terminator

monitor_field_declaration ::=
    const_field_declaration
  | action_handle_declaration
  | monitor_handle_declaration

monitor_activity_declaration ::=
  activity { { monitor_activity_stmt } }

monitor_activity_stmt ::=
    [ label_identifier : ] labeled_monitor_activity_stmt
  | monitor_activity_action_traversal_stmt
  | monitor_activity_monitor_traversal_stmt
  | action_handle_declaration
  | monitor_handle_declaration
  | monitor_activity_constraint_stmt
  | annotation_usage monitor_activity_stmt
  | stmt_terminator

labeled_monitor_activity_stmt ::=
    monitor_activity_sequence_block_stmt
  | monitor_activity_concat_stmt
  | monitor_activity_eventually_stmt
  | monitor_activity_overlap_stmt
  | monitor_activity_schedule_stmt
  | monitor_activity_select_stmt
  | activity_super_stmt
```

Copyright © 2026 Accellera. All rights reserved.

```
monitor_handle_declaration ::= monitor_type_identifier
    monitor_instantiation ;

monitor_instantiation ::=
  monitor_identifier { array_dim }
    { , monitor_identifier { array_dim }}

monitor_activity_action_traversal_stmt ::= activity_action_traversal_stmt

monitor_activity_sequence_block_stmt ::= [ sequence ] {{monitor_activity_stmt}}

monitor_activity_concat_stmt ::= concat {{monitor_activity_stmt}}

monitor_activity_eventually_stmt ::= eventually monitor_activity_stmt ;

monitor_activity_overlap_stmt ::= overlap {{monitor_activity_stmt}}

monitor_activity_select_stmt ::= select {monitor_activity_stmt
monitor_activity_stmt { monitor_activity_stmt }}

monitor_activity_schedule_stmt ::= schedule {{monitor_activity_stmt}}

monitor_activity_monitor_traversal_stmt ::=
    monitor_handle_traversal_stmt
    | monitor_type_traversal_stmt

monitor_handle_traversal_stmt ::=
     monitor_identifier { [ expression ] } [ action_initializer_list ]
       monitor_activity_inline_constraints_or_empty

monitor_type_traversal_stmt ::=
  [ label_identifier : ] do monitor_type_identifier [action_initializer_list]
      monitor_activity_inline_constraints_or_empty

monitor_activity_inline_constraints_or_empty ::=
     with monitor_constraint_set
  | ;

monitor_activity_constraint_stmt ::= constraint monitor_constraint_set

monitor_constraint_declaration ::=
     constraint monitor_constraint_set
  | constraint identifier monitor_constraint_block

monitor_constraint_set ::=
     monitor_constraint_body_item
  | monitor_constraint_block

monitor_constraint_block ::= { { monitor_constraint_body_item } }

monitor_constraint_body_item ::=
     expression_constraint_item
  | foreach_constraint_item
  | forall_constraint_item
  | if_constraint_item
  | implication_constraint_item
  | unique_constraint_item
```

Copyright © 2026 Accellera. All rights reserved.

```
  | constraint_body_compile_if
  | annotation_usage monitor_constraint_body_item
  | stmt_terminator
```

B.17 Conditional compilation

```
package_body_compile_if ::= compile if ( constant_expression )
    package_body_compile_if_item [ else package_body_compile_if_item ]

monitor_body_compile_if ::= compile if ( constant_expression )
  monitor_body_compile_if_item [ else monitor_body_compile_if_item ]

action_body_compile_if ::= compile if ( constant_expression )
    action_body_compile_if_item [ else action_body_compile_if_item ]

component_body_compile_if ::= compile if ( constant_expression )
    component_body_compile_if_item [ else component_body_compile_if_item ]

struct_body_compile_if ::= compile if ( constant_expression )
    struct_body_compile_if_item [ else struct_body_compile_if_item ]

procedural_compile_if ::= compile if ( constant_expression )
    procedural_compile_if_stmt [ else procedural_compile_if_stmt ]

constraint_body_compile_if ::= compile if ( constant_expression )
    constraint_body_compile_if_item [ else constraint_body_compile_if_item ]

covergroup_body_compile_if ::= compile if ( constant_expression )
    covergroup_body_compile_if_item [ else covergroup_body_compile_if_item ]

override_compile_if ::= compile if ( constant_expression )
    override_compile_if_stmt [ else override_compile_if_stmt ]

package_body_compile_if_item ::= { { package_body_item } }

action_body_compile_if_item ::= { { action_body_item } }

monitor_body_compile_if_item ::= { { monitor_body_item } }

component_body_compile_if_item ::= { { component_body_item } }

struct_body_compile_if_item ::= { { struct_body_item } }

procedural_compile_if_stmt ::= { { procedural_stmt } }

constraint_body_compile_if_item ::= { { constraint_body_item } }

covergroup_body_compile_if_item ::= { { covergroup_body_item } }

override_compile_if_stmt ::= { { override_stmt } }

compile_has_expr ::= compile has ( static_ref_path )

compile_assert_stmt ::=
```

Copyright © 2026 Accellera. All rights reserved.

```
compile assert ( constant_expression [ , string_literal ] ) ;
```

B.18 Expressions

```
constant_expression ::= expression

expression ::=
      primary
    | unary_operator primary
    | expression binary_operator expression
    | conditional_expression
    | in_expression

unary_operator ::= - | ! | ~ | & | | | ^

binary_operator ::=
      * | / | % | + | - | << | >> | == | != | < | <= | > | >= | || | && | |
    | ^ | & | **

assign_op ::= = | += | -= | <<= | >>= | |= | &=

conditional_expression ::= cond_predicate ? expression : expression

cond_predicate ::= expression

in_expression ::=
      expression in [ open_range_list ]
    | expression in collection_expression [ [ array_slice ] ]

open_range_list ::= open_range_value { , open_range_value }

open_range_value ::= expression [ .. expression ]

collection_expression ::= expression

array_slice ::=
      expression .. expression
    | expression ..
    | .. expression

primary ::=
      number
    | aggregate_literal
    | bool_literal
    | string_literal
    | null_ref
    | paren_expr
    | cast_expression
    | ref_path
    | compile_has_expr

paren_expr ::= ( expression )

cast_expression ::= ( casting_type ) expression

ref_path ::=
```

Copyright © 2026 Accellera. All rights reserved.

```
      static_ref_path [ . hierarchical_id ] [ slice ]
    | [ super . ] hierarchical_id [ slice ]
slice ::= bit_slice | string_slice

static_ref_path ::= [ :: ] { type_identifier_elem :: } member_path_elem

bit_slice ::= [ constant_expression : constant_expression ]

string_slice ::=
     expression [ .. expression ]
   | expression ..
   | .. expression

function_call ::=
      super . function_ref_path
    | [ :: ] { type_identifier_elem :: } function_ref_path

function_ref_path ::= { member_path_elem . } identifier function_parameter_list

symbol_call ::= symbol_identifier function_parameter_list ;

function_parameter_list ::= ( [ expression { , expression } ] )
```

B.19 Identifiers

```
identifier ::=
      ID
    | ESCAPED_ID

hierarchical_id_list ::= hierarchical_id { , hierarchical_id }

hierarchical_id ::= member_path_elem { . member_path_elem }

member_path_elem ::= identifier [ function_parameter_list ] { [ expression ] }

action_identifier ::= identifier

action_handle_identifier ::= identifier

component_identifier ::= identifier

covercross_identifier ::= identifier

covergroup_identifier ::= identifier

coverpoint_identifier ::= identifier

enum_identifier ::= identifier

function_identifier ::= identifier

import_class_identifier ::= identifier

index_identifier ::= identifier

iterator_identifier ::= identifier
```

Copyright © 2026 Accellera. All rights reserved.

```
label_identifier ::= identifier

language_identifier ::= identifier

monitor_identifier ::= identifier

package_identifier ::= identifier

struct_identifier ::= identifier

symbol_identifier ::= identifier

type_identifier ::= [ :: ] type_identifier_elem { :: type_identifier_elem }

type_identifier_elem ::= identifier [ template_param_value_list ]

action_type_identifier ::= type_identifier

buffer_type_identifier ::= type_identifier

component_type_identifier ::= type_identifier

covergroup_type_identifier ::= type_identifier

enum_type_identifier ::= type_identifier

monitor_type_identifier ::= type_identifier

resource_type_identifier ::= type_identifier

state_type_identifier ::= type_identifier

stream_type_identifier ::= type_identifier

entity_type_identifier ::=
      action_type_identifier
    | monitor_type_identifier
    | component_type_identifier
    | flow_object_type
    | resource_object_type
```

B.20 Numbers and literals

```
number ::=
      integer_number
    | floating_point_number

integer_number ::=
      bin_number
    | oct_number
    | dec_number
    | hex_number
    | based_bin_number
    | based_oct_number
    | based_dec_number
    | based_hex_number
```

Copyright © 2026 Accellera. All rights reserved.

```
bin_digit ::= [0-1]

oct_digit ::= [0-7]

dec_digit ::= [0-9]

hex_digit ::= [0-9] | [a-f] | [A-F]

bin_number ::= 0[b|B] bin_digit { bin_digit | _ }

oct_number ::= 0 { oct_digit | _ }

dec_number ::= [1-9] { dec_digit | _ }

hex_number ::= 0[x|X] hex_digit { hex_digit | _ }

BASED_BIN_LITERAL ::= '[s|S]b|B bin_digit { bin_digit | _ }

BASED_OCT_LITERAL ::= '[s|S]o|O oct_digit { oct_digit | _ }

BASED_DEC_LITERAL ::= '[s|S]d|D dec_digit { dec_digit | _ }

BASED_HEX_LITERAL ::= '[s|S]h|H hex_digit { hex_digit | _ }

based_bin_number ::= [ dec_number ] BASED_BIN_LITERAL

based_oct_number ::= [ dec_number ] BASED_OCT_LITERAL

based_dec_number ::= [ dec_number ] BASED_DEC_LITERAL

based_hex_number ::= [ dec_number ] BASED_HEX_LITERAL

floating_point_number ::=
      floating_point_dec_number
    | floating_point_sci_number

unsigned_number ::= dec_digit { dec_digit | _ }

floating_point_dec_number ::= unsigned_number . unsigned_number

floating_point_sci_number ::=
      unsigned_number [ . unsigned_number ] exp [ sign ] unsigned_number

exp ::= e | E

sign ::= + | -

aggregate_literal ::=
      empty_aggregate_literal
    | value_list_literal
    | map_literal
    | struct_literal

empty_aggregate_literal ::= { }

value_list_literal ::= { expression { , expression } }
```

Copyright © 2026 Accellera. All rights reserved.

```
map_literal ::= { map_literal_item { , map_literal_item } }

map_literal_item ::= expression : expression

struct_literal ::= { struct_literal_item { , struct_literal_item } }

struct_literal_item ::= . identifier = expression

bool_literal ::=
      true
    | false

null_ref ::= null
```

B.21 Additional lexical conventions

```
SL_COMMENT ::= //{any_ASCII_character_except_newline}\n

ML_COMMENT ::= /*{any_ASCII_character}*/

string_literal ::=
      QUOTED_STRING
    | TRIPLE_QUOTED_STRING

QUOTED_STRING ::= " { unescaped_character | escaped_character } "

unescaped_character ::= any_printable_ASCII_character

escaped_character ::= \('|"|?|\|a|b|f|n|r|t|v|[0-7][0-7][0-7])

TRIPLE_QUOTED_STRING ::= """{any_ASCII_character}"""

filename_string ::= string_literal

ID ::= [a-z]|[A-Z]|_ {[a-z]|[A-Z]|_|[0-9]}

ESCAPED_ID ::= \{any_printable_ASCII_character_except_whitespace} whitespace

whitespace ::= space | tab | newline | end_of_file
```

Copyright © 2026 Accellera. All rights reserved.

Annex C
(normative)
Core library package

This annex contains the contents of the built-in core library packages **std_pkg**, **executor_pkg**, **addr_reg_pkg**, and **sync_pkg** described in Clause 21. If there is a conflict between core library package contents shown anywhere in this standard and the material in this annex, the material shown in this annex shall take precedence.

C.1 Package std_pkg

```
package std_pkg {

    enum endianness_e {LITTLE_ENDIAN, BIG_ENDIAN};

    struct packed_s<endianness_e e = LITTLE_ENDIAN> {};

    struct sizeof_s<type T> {
        static const int nbytes = /* implementation-specific */;
        static const int nbits = /* implementation-specific */;
    };


    // Functions available on solve platform only
    solve pure function string format(string format_str, type... args);
    solve function void print(string format_str, type... args);

    enum message_verbosity_e {NONE, LOW, MEDIUM, HIGH, FULL};

    function void message
        (message_verbosity_e vrb_level, string format_str, type... args);

    typedef chandle file_handle_t;
    static const file_handle_t nullfilehandle = /* implementation-specific */;

    enum file_option_e {TRUNCATE, APPEND, READ};

    // Functions available on solve platform only
    solve function file_handle_t file_open(string filename, file_option_e opt);
    solve function void file_close(file_handle_t file_handle);
    solve function bool file_exists(string filename);

    solve function void file_write
        (file_handle_t file_handle, string format_str, type... args);
    solve function string file_read(file_handle_t file_handle, int size = -1);

    solve function void file_write_lines
       (string filename, list<string> lines, file_option_e opt);
    solve function list<string> file_read_lines(string filename);

    function void error(string format_str, type... args);
    function void fatal(int status, string format_str, type... args);
```

Copyright © 2026 Accellera. All rights reserved.

```
    // random functions
    function bit[32] urandom();
    function bit[32] urandom_range(bit[32] min, bit[32] max);


    // Floating-point Storage Types
    struct float_base_s <int Wm, int We, endianness_e E=LITTLE_ENDIAN> :
       packed_s<E> {
           rand bit[Wm] mantissa;
           rand bit[We] exponent;
           rand bit     sign;
       }

    // Pre-defined storage types to match computation types
    typedef float_base_s<23, 8> float32_s;
    typedef float_base_s<52,11> float64_s;

    // Floating-point Functions
    pure function float64 log(float64 x);
    pure function float64 log10(float64 x);
    pure function float64 exp(float64 x);
    pure function float64 sqrt(float64 x);
    pure function float64 pow(float64 x, float64 y);
    pure function float64 round(float64 x);
    pure function float64 floor(float64 x);
    pure function float64 ceil(float64 x);
    pure function float64 sin(float64 x);
    pure function float64 cos(float64 x);
    pure function float64 tan(float64 x);
    pure function float64 asin(float64 x);
    pure function float64 acos(float64 x);
    pure function float64 atan(float64 x);
    pure function float64 atan2(float64 y, float64 x);
    pure function float64 hypot(float64 x, float64 y);
    pure function float64 sinh(float64 x);
    pure function float64 cosh(float64 x);
    pure function float64 tanh(float64 x);
    pure function float64 asinh(float64 x);
    pure function float64 acosh(float64 x);
    pure function float64 atanh(float64 x);

    pure function bit[52] float_mantissa(float64 fv);
    pure function bit[11] float_exponent(float64 fv);
    pure function bit float_sign(float64 fv);
    pure function float64 to_float(bit[52] mantissa, bit[11] exp, bit sign);

  annotation code_doc {
    string text;
  }

  annotation doc {
    string text;
  }
}
```

Copyright © 2026 Accellera. All rights reserved.

C.2 Package executor_pkg

```
package executor_pkg {

   struct executor_trait_s {};

   struct empty_executor_trait_s : executor_trait_s {};
   enum target_language_e {C, CPP, SV};
   component target_execution_unit_c {
     solve function void set_filename(string filename);
     solve function string get_filename();
     solve function void set_target_language(target_language_e lang);
     solve function target_language_e get_target_language();
   }

  component executor_base_c {
    target function chandle get_context();
      solve function void set_target_execution_unit(
          ref target_execution_unit_c target_execution_unit);
      solve function ref target_execution_unit_c get_target_execution_unit();
  }

   component executor_c
                <struct TRAIT : executor_trait_s = empty_executor_trait_s>
                : executor_base_c {
     TRAIT trait;
   };

   component executor_group_c
                <struct TRAIT : executor_trait_s = empty_executor_trait_s> {
                solve function void add_executor(ref executor_c<TRAIT> exe);
   };

   struct executor_claim_s
                <struct TRAIT : executor_trait_s = empty_executor_trait_s> {
      rand TRAIT trait;
   };

   function ref executor_base_c executor();
   function void set_executor(ref executor_base_c xtr);
}
```

C.3 Package addr_reg_pkg

```
package addr_reg_pkg {
   import std_pkg::* ;
   import executor_pkg::* ;

   component addr_space_base_c {};

   struct addr_trait_s {};

   struct empty_addr_trait_s : addr_trait_s {};

   typedef chandle addr_handle_t;

   component contiguous_addr_space_c
```

Copyright © 2026 Accellera. All rights reserved.

```
                    <struct TRAIT : addr_trait_s = empty_addr_trait_s>
                    : addr_space_base_c {

   solve function addr_handle_t add_region(addr_region_s <TRAIT> r);
   solve function addr_handle_t
                    add_nonallocatable_region (addr_region_s <> r);

   bool byte_addressable = true;
};

component transparent_addr_space_c
                    <struct TRAIT: addr_trait_s = empty_addr_trait_s>
                    : contiguous_addr_space_c<TRAIT> {};
 component addr_space_group_c  {
    function void add_addr_space(ref addr_space_base_c address_space);
 }

struct addr_region_base_s {
   bit[64] size;
   string  tag;

};

struct addr_region_s <struct TRAIT : addr_trait_s = empty_addr_trait_s>
                     : addr_region_base_s {
   TRAIT trait;
};

 struct transparent_addr_region_s
                    <struct TRAIT : addr_trait_s = empty_addr_trait_s>
                    : addr_region_s<TRAIT> {
   bit[64] addr;
};

 enum alloc_access_mode_e {EXCLUSIVE, SHARED};
 enum alloc_share_mode_e {RANDOM, OVERLAP, SAME_ADDRESS};
 enum alloc_skip_mode_e {DONT_SKIP, SKIP};

 struct alloc_base_mode_s {
    rand alloc_access_mode_e access;
    constraint default access == EXCLUSIVE;
    rand alloc_share_mode_e sharing;
    constraint default sharing == RANDOM;
    rand alloc_skip_mode_e skipping;
    constraint default skipping == DONT_SKIP;
    rand int tag;
    constraint default tag == -1;
 };

 struct addr_claim_base_s {
    rand bit[64] size;
    rand bool permanent;
    constraint default permanent == false;
 };
```

Copyright © 2026 Accellera. All rights reserved.

```
struct addr_claim_s
              <struct TRAIT : addr_trait_s = empty_addr_trait_s,
               struct ALLOC_MODE : alloc_base_mode_s = alloc_base_mode_s>
              :addr_claim_base_s {
  rand TRAIT trait;
  rand bit[64] in [64'd2**0, 64'd2**1, 64'd2**2, 64'd2**3 , 64'd2**4 ,
     64'd2**5 , 64'd2**6 , 64'd2**7 , 64'd2**8 , 64'd2**9 , 64'd2**10,
     64'd2**11, 64'd2**12, 64'd2**13, 64'd2**14, 64'd2**15, 64'd2**16,
     64'd2**17, 64'd2**18, 64'd2**19, 64'd2**20, 64'd2**21, 64'd2**22,
     64'd2**23, 64'd2**24, 64'd2**25, 64'd2**26, 64'd2**27, 64'd2**28,
     64'd2**29, 64'd2**30, 64'd2**31, 64'd2**32, 64'd2**33, 64'd2**34,
     64'd2**35, 64'd2**36, 64'd2**37, 64'd2**38, 64'd2**39, 64'd2**40,
     64'd2**41, 64'd2**42, 64'd2**43, 64'd2**44, 64'd2**45, 64'd2**46,
     64'd2**47, 64'd2**48, 64'd2**49, 64'd2**50, 64'd2**51, 64'd2**52,
     64'd2**53, 64'd2**54, 64'd2**55, 64'd2**56, 64'd2**57, 64'd2**58,
     64'd2**59, 64'd2**60, 64'd2**61, 64'd2**62, 64'd2**63] alignment;
  rand ALLOC_MODE alloc_mode;
};

struct transparent_addr_claim_s
 <struct TRAIT : addr_trait_s = empty_addr_trait_s,
   struct ALLOC_MODE : alloc_base_mode_s = alloc_base_mode_s > :
   addr_claim_s<TRAIT, ALLOC_MODE> {
  rand bit[64] addr;
}

const addr_handle_t nullhandle = /* implementation-specific */;

struct sized_addr_handle_s < int SZ, // in bits
                             int lsb = 0,
                             endianness_e e = LITTLE_ENDIAN >
                           : packed_s<e> {
   addr_handle_t hndl;
};


function addr_handle_t make_handle_from_claim  (addr_claim_base_s claim,
                                                bit[64] offset = 0,
                                                bool sub = false);

function addr_handle_t make_handle_from_handle (addr_handle_t handle,
                                                bit[64] offset,
                                                bool sub = false);

struct mem_access_desc_s {}
target function bit[64] addr_value(addr_handle_t hndl,
                           const mem_access_desc_s mem_access_desc = {});
solve function bit[64] addr_value_solve(addr_handle_t hndl,
                           const mem_access_desc_s mem_access_desc = {});
solve function bool addr_value_abs(addr_handle_t hndl,
                           const mem_access_desc_s mem_access_desc = {});

function string get_tag(addr_handle_t hndl);

target function bit[8] read8(addr_handle_t hndl,
                           const mem_access_desc_s mem_access_desc = {});
target function bit[16] read16(addr_handle_t hndl,
                           const mem_access_desc_s mem_access_desc = {});
target function bit[32] read32(addr_handle_t hndl,
```

Copyright © 2026 Accellera. All rights reserved.

```
                                const mem_access_desc_s mem_access_desc = {});
   target function bit[64] read64(addr_handle_t hndl,
                                 const mem_access_desc_s mem_access_desc = {});

   target function void write8 (addr_handle_t hndl, bit[8] data,
                                const mem_access_desc_s mem_access_desc = {});
   target function void write16(addr_handle_t hndl, bit[16] data,
                                const mem_access_desc_s mem_access_desc = {});
   target function void write32(addr_handle_t hndl, bit[32] data,
                                const mem_access_desc_s mem_access_desc = {});
   target function void write64(addr_handle_t hndl, bit[64] data,
                                const mem_access_desc_s mem_access_desc = {});

   target function void read_bytes (addr_handle_t hndl, list<bit[8]> data,
                                    int size,
                                const mem_access_desc_s mem_access_desc = {});
   target function void write_bytes(addr_handle_t hndl,
                                    list<bit[8]> data,
                                 const mem_access_desc_s mem_access_desc = {});

   target function void read_struct (addr_handle_t hndl, struct packed_struct,
                                const mem_access_desc_s mem_access_desc = {});
   target function void write_struct(addr_handle_t hndl,
                                     struct packed_struct,
                                 const mem_access_desc_s mem_access_desc = {});

extend component executor_base_c {
   target function bit[64] addr_value(addr_handle_t hndl,
                                 const mem_access_desc_s mem_access_desc = {});
     solve function bit[64] addr_value_solve(addr_handle_t hndl,
                                const mem_access_desc_s mem_access_desc = {});
     solve function bool addr_value_abs(addr_handle_t hndl,
                                 const mem_access_desc_s mem_access_desc = {});
   target function bit[8] read8(addr_handle_t hndl,
                                 const mem_access_desc_s mem_access_desc = {});
   target function bit[16] read16(addr_handle_t hndl,
                                 const mem_access_desc_s mem_access_desc = {});
   target function bit[32] read32(addr_handle_t hndl,
                                 const mem_access_desc_s mem_access_desc = {});
   target function bit[64] read64(addr_handle_t hndl,
                                 const mem_access_desc_s mem_access_desc = {});
   target function void write8 (addr_handle_t hndl, bit[8] data,
                                const mem_access_desc_s mem_access_desc = {});
   target function void write16(addr_handle_t hndl, bit[16] data,
                                const mem_access_desc_s mem_access_desc = {});
   target function void write32(addr_handle_t hndl, bit[32] data,
                                const mem_access_desc_s mem_access_desc = {});
   target function void write64(addr_handle_t hndl, bit[64] data,
                                const mem_access_desc_s mem_access_desc = {});
   target function void read_bytes(addr_handle_t hndl, list<bit[8]> data,
                                   int size,
                                const mem_access_desc_s mem_access_desc = {});
   target function void write_bytes(addr_handle_t hndl, list<bit[8]> data,
                                const mem_access_desc_s mem_access_desc = {});
};

   pure component reg_base_c {
     function addr_handle_t get_handle();
   }
```

Copyright © 2026 Accellera. All rights reserved.

```
    pure component reg_sized_c<int SZ> : reg_base_c {
      target function bit[SZ] read_val
                             (const mem_access_desc_s mem_access_desc = {});
      target function void write_val(bit[SZ] r,
                             const mem_access_desc_s mem_access_desc = {});
      target function void write_val_masked
                                  (bit[SZ] mask, bit[SZ] val,
                             const mem_access_desc_s mem_access_desc = {});
      target function void write_field
                                  (string name, bit[SZ] val,
                             const mem_access_desc_s mem_access_desc = {});
      target function void write_fields
                                  (list<string> names, list<bit[SZ]> vals,
                             const mem_access_desc_s mem_access_desc = {});
    }

    enum reg_access {READWRITE, READONLY, WRITEONLY};

    pure component reg_c <type R,
    reg_access ACC = READWRITE,
      int SZ = (8*sizeof_s<R>::nbytes)>: reg_sized_c<SZ>{
        target function R read(const mem_access_desc_s mem_access_desc = {});
        target function void write(R r,
                             const mem_access_desc_s mem_access_desc = {});
        target function void write_masked
                                (R mask, R val,
                             const mem_access_desc_s mem_access_desc = {});
    };

    struct node_s {
       string name;
       int    index;
    };

    solve function void use_symbolic_reg_names(ref reg_group_c grp,
                                               bool enable);
    pure component reg_group_c {
       pure function bit[64] get_offset_of_instance(string name);
       pure function bit[64] get_offset_of_instance_array
                                                 (string name, int index);
       pure function bit[64] get_offset_of_path(list<node_s> path);
       solve function void set_handle(addr_handle_t addr);
       function addr_handle_t get_handle();

       solve pure function string get_mnemonic_of_instance(string name);
       solve pure function string get_mnemonic_of_instance_array(string name,
                                                          int index);
       solve pure function string get_mnemonic_of_path(list<node_s> path);
       solve function void        set_mnemonic(string prefix);
};
```

C.4 Package sync_pkg

```
package sync_pkg {
  component channel_c<type T, int DEPTH=1> {
    target function T get();
```

Copyright © 2026 Accellera. All rights reserved.

```
    target function void put(T t);
    target function bool try_get(output T t);
    target function bool try_put(T t);
  }
}
```

Copyright © 2026 Accellera. All rights reserved.

Annex D
(normative)
Foreign language bindings

D.1 Function prototype mapping

Let *f* be a function declared under hierarchical path *H* in PSS with type signature as below (with D_x as the direction, T_x as the type and p_x as the parameter name):

$$f(D_0\ T_0\ p_0,\ D_1\ T_1\ p_1,\ \ldots,\ D_n\ T_n\ p_n);$$

When *f* is bound to a foreign language API (see 20.4), it is mapped to the following function in the target language:

$$H'::f'(T'_0\ p_0,\ T'_1\ p_1,\ \ldots,\ T'_n\ p_n);$$

If the foreign language supports parameter directions, their directions are the same as in PSS.

NOTE—See D.5 for exceptions when mapping PSS functions to SystemVerilog tasks.

Each parameter in the PSS function is mapped to a corresponding parameter in the mapped function. The details of function name and data type binding are covered further below.

D.1.1 Export functions

The first parameter of the foreign-language realization of a PSS exported function establishes the executor context in which the function executes. The type of this context parameter is foreign-language specific. Consequently, the signature of the foreign-language realization of a PSS exported function is as follows, where `CT` is the foreign-language-specific context-handle type:

```
H'::f'(CT p0, T'0 p1, T'1 p2, . . ., T'n pn+1);
```

The following additional restrictions apply to the parameters of export functions:

a) Parameters of type **array**, **list**, **set**, and map-type are not supported.
b) Parameters of struct type shall be declared as **const**.
c) All parameters are treated as **input** for the purposes of the mapping tables in this annex.

D.2 Data type mapping

PSS specifies data type bindings to C/C++ and SystemVerilog. The data type binding rules apply only to parameter and return types referenced (directly or indirectly) in the declaration of functions in PSS that are bound to foreign language APIs (see 20.4). The allowed types are specified in 20.4.1.1, namely:

— Primitive types: **bit** or **int** (width no more than 64 bits), **bool**, **string**, **chandle**.
— User-defined types: **enum** and **struct**, excluding packed structs (see 21.11.1) and excluding flow/resource objects. Fields of **struct**s shall be of these allowed types (recursively).
— Fixed-size arrays of these types.

The type binding is specified for parameter and return types.

Copyright © 2026 Accellera. All rights reserved.

D.3 C language bindings

D.3.1 Function names

PSS implementations shall support mapping a PSS function name to an identical function name in C, ignoring the hierarchical path in PSS. PSS implementations may define additional mapping schemes for function names.

D.3.2 Primitive types

The mapping between the PSS primitive types and C types is specified in Table D1.

Table D.1—Mapping PSS primitive types and C types

PSS type	C input type	C output/inout type	C return type
string	**const char ***	**char ****	**char ***
bool	**unsigned int**	**unsigned int ***	**unsigned int**
chandle	**const void ***	**void ****	**void ***
bit (1-8-bit domain)	**unsigned char**	**unsigned char ***	**unsigned char**
bit (9-16-bit domain)	**unsigned short**	**unsigned short ***	**unsigned short**
bit (17-32-bit domain)	**unsigned int**	**unsigned int ***	**unsigned int**
bit (33-64-bit domain)	**unsigned long long**	**unsigned long long ***	**unsigned long long**
int (1-8-bit domain)	**char**	**char ***	**char**
int (9-16-bit domain)	**short**	**short ***	**short**
int (17-32-bit domain)	**int**	**int ***	**int**
int (33-64-bit domain)	**long long**	**long long ***	**long long**
float32	**float**	**float ***	**float**
float64	**double**	**double ***	**double**

Where pointers are used, the callee shall not allocate or de-allocate the memory region referenced by the pointer. Further, for non-void pointers, the callee shall assume that the memory location is valid only for the duration of the function execution, and shall not retain a reference to the parameter after the function call returns. For **string**s and **chandle**s, in the case of **inout**/**output** directions, the callee may return a pointer to storage it owns.

D.3.3 Arrays

Fixed-sized arrays are mapped to fixed-size arrays in C for function arguments. Mapping PSS fixed-sized arrays to C is not supported for function return types.

Copyright © 2026 Accellera. All rights reserved.

D.3.4 Structs

D.3.4.1 Name mapping

The mapping between a PSS **struct** type (T_{PSS}) defined in a hierarchical path H and a C type (T_C) is shown in Table D2.

Table D.2—Mapping PSS struct types and C types

PSS type	C input type	C output/inout type	C return type
$H::T_{PSS}$	const T_C *	T_C *	T_C

In the general case, the name of the type in C (T_C), is derived from the PSS type name (T_{PSS}) and its hierarchical path (H). A PSS implementation shall support the name mapping scheme where the name of the C type is identical to the PSS type (ignoring the hierarchical path), i.e., $T_C == T_{PSS}$. A PSS implementation may support additional name mapping schemes.

D.3.4.2 Field mapping

Each PSS **struct** field is mapped to a corresponding field in C of the corresponding type and name in the same order. If the field type is itself a user-defined type (e.g., **struct** or **enum**), the mapping of the field entails the corresponding mapping of the type (recursively). For primitive types, the field is mapped as shown in Table D3.

Table D.3—Mapping PSS struct field primitive types and C types

PSS field type	C field type
string	**char ***
bool	**unsigned int**
chandle	**void ***
bit (1-8-bit domain)	**unsigned char**
bit (9-16-bit domain)	**unsigned short**
bit (17-32-bit domain)	**unsigned int**
bit (33-64-bit domain)	**unsigned long long**
int (1-8-bit domain)	**char**
int (9-16-bit domain)	**short**
int (17-32-bit domain)	**int**
int (33-64-bit domain)	**long long**
float32	**float**
float64	**double**

Copyright © 2026 Accellera. All rights reserved.

A PSS implementation shall support mapping a PSS field name to an identical field name in C. A PSS implementation may support additional mapping schemes for field names.

Since the C language does not support type inheritance, if the PSS **struct** T_{PSS} derives from a PSS base type, then the fields of that base type are mapped directly into the mapped type T_C. The code listing below shows an example of **struct** type mapping in C.

```
// PSS code

struct base_s {
  string f0;
};

struct sub_s {
  int[8] f1 = 2;
  string f2;
};


struct my_struct_s : base_s {
  sub_s f3;
  bit[16] f4;
};

my_struct_s function foo
    (input my_struct_s x,
    output my_struct_s y);
```

```
// C code

struct sub_s {
  char f1;
  char *f2;
};

struct my_struct_s {
  char *f0;
  sub_s f3;
  unsigned short f4;
};

my_struct_s foo
    (const my_struct_s *x,
           my_struct_s *y);
```

Example D.1—PSS struct mapping into C

Only the field name, its type and the position of the field inside a **struct** is relevant for mapping to the C type. Other field properties (such as initial value) and **struct** properties (such as constraints) are ignored.

D.3.4.3 Other mapping aspects

Tools may automatically generate C definitions for the required types, given PSS source code. Or, tools may utilize existing C declarations of the types. Regardless of whether these definitions are automatically generated or obtained in another way, PSS test generation tools may assume that these definitions are operative in the compilation of the C user implementation of the imported functions.

Note that the C declaration of a **struct** data type may have additional fields that are not reflected in the PSS type declaration. Such fields shall be declared after those reflecting the fields in the PSS type declaration. A PSS implementation may not assume that the C struct is size-compatible to the PSS **struct** type.

Copyright © 2026 Accellera. All rights reserved.

D.3.5 Enumeration types

A PSS enumeration type *E* is mapped to C as a plain integer type *N* as follows:

Table D.4—Mapping PSS enum types and C types

PSS type	C input type	C output/inout type	C return type
E	***N***	***N****	***N***

where *N* is:

a) If E has a base type: the mapping for the base type, according to D.3.2.

b) Otherwise, one of: **char**, **short**, **int**, or **long long**, according to the smallest type that includes the values of all the enum items in its domain.

A PSS implementation will pass the value of the enumeration as an argument in the generated call to the function. These values can be either explicitly user-defined or assigned by a PSS implementation.

D.3.6 Export-function context handle

In C, a void pointer is used as the type for export-function context handles.

Example D.2 shows the definition of a PSS function that is exported. Example D.3 shows the C prototype for that function.

```
function void exp_func(int val) {
}
export target function exp_func;
```

Example D.2—PSS-defined exported function

```
void exp_func(void *context, int val);
```

Example D.3—C prototype for exported function

D.4 C++ language bindings

D.4.1 Function name mapping and namespaces

Generally, PSS user-defined types correspond to C++ types with identical names. In PSS, **packages** and **components** constitute namespaces for types declared in their scopes. The C++ type definition corresponding to a PSS type declared in a **package** or **component** scope shall be inside the namespace statement scope having the same name as the PSS **component**/**package**. Consequently, both the unqualified and qualified names of the C++ mapped type are the same as in PSS.

Copyright © 2026 Accellera. All rights reserved.

PSS implementations shall support mapping a PSS function name to an identical function name in C++, in the same namespace hierarchical path. PSS implementations may define additional mapping schemes for function names.

D.4.2 Primitive types

a) C++ type mapping for primitive numeric types is the same as that for C.

b) A PSS **bool** is a C++ **`bool`** and the values: **`false`**, **`true`** are mapped respectively from PSS to their C++ equivalents.

c) C++ mapping of a PSS **string** is **`std::string`** (typedef-ed by the Standard Template Library (STL) to **`std::basic_string<char>`** with default template parameters).

Table D5 provides the mapping between PSS primitive types and C++ types. Note that **string** is passed as a reference.

Table D.5—Mapping PSS primitive types and C++ types

PSS type	C++ input type	C++ output/inout type	C++ return type
string	**const std::string &**	**std::string &**	**std::string**
bool	**bool**	**bool ***	**bool**
chandle	**const void ***	**void ****	**void ***
bit (1-8-bit domain)	**unsigned char**	**unsigned char ***	**unsigned char**
bit (9-16-bit domain)	**unsigned short**	**unsigned short ***	**unsigned short**
bit (17-32-bit domain)	**unsigned int**	**unsigned int ***	**unsigned int**
bit (33-64-bit domain)	**unsigned long long**	**unsigned long long ***	**unsigned long long**
int (1-8-bit domain)	**char**	**char ***	**char**
int (9-16-bit domain)	**short**	**short ***	**short**
int (17-32-bit domain)	**int**	**int ***	**int**
int (33-64-bit domain)	**long long**	**long long ***	**long long**
float32	**float**	**float ***	**float**
float64	**double**	**double ***	**double**

D.4.3 Arrays

The C++ mapping of a PSS array is **`std::vector`** of the C++ mapping of the respective element type (using the default allocator class). Fixed-sized arrays in PSS are mapped to the corresponding STL vector class, just like arrays of an unspecified size. However, if modified, they are resized to the original size upon return, filling the default values of the respective element type as needed.

Copyright © 2026 Accellera. All rights reserved.

D.4.4 Structs

D.4.4.1 Name mapping

The mapping between a PSS **struct** type (T_{PSS}) and a C++ type (T_{CPP}) is shown in Table D6.

Table D.6—Mapping PSS struct types and C++ types

PSS type	C++_ input type	C++ output/inout type	C++ return type
T_{PSS}	const T_{CPP} &	T_{CPP} &	T_{CPP}

PSS **struct** types are mapped to C++ structs, along with their field structure and inherited base type, if specified.

The base type declaration of the **struct**, if any, is mapped to the (public) base struct type declaration in C++ and entails the mapping of its base type (recursively).

D.4.4.2 Field mapping

Each PSS field is mapped to a corresponding (public, non-static) field in C++ of the corresponding type and in the same order. If the field type is itself a user-defined type (**struct** or **enum**), the mapping of the field entails the corresponding mapping of the type (recursively).

A PSS implementation shall support mapping a PSS field name to an identical field name in C++. A PSS implementation may support additional mapping schemes for field names.

For example, given the following imported function definitions:

```
function void foo(derived_s d);
import solve CPP function foo;
```

with the corresponding PSS definitions:

```
struct base_s {
   int in [0..99] f1;
};
struct sub_s {
   string f2;
};
struct derived_s : base_s {
   sub_s f3;
   bit[16] f4[4];
};
```

mapping type `derived_s` to C++ involves the following definitions:

```
struct base_s {
   int f1;
};
struct sub_s {
   std::string f2;
};
```

Copyright © 2026 Accellera. All rights reserved.

```
struct derived_s : base_s {
   sub_s f3;
   std::vector<unsigned short> f4;
};
```

Nested **structs** in PSS are instantiated directly under the containing **struct**, that is, they have value semantics. Mapped **struct** types have no member functions and, in particular, are confined to the default constructor and implicit copy constructor.

Mapping a **struct** type does not entail the mapping of any of its subtypes. However, **struct** instances are passed according to the type of the actual parameter expression used in an **import function** call. Therefore, the ultimate set of C++ mapped types for a given PSS model depends on its function calls, not just the function prototypes.

D.4.4.3 Other mapping aspects

In the case of **`output`** and **`inout`** aggregate parameters, if a different memory representation is used for the PSS tool vs. C++, the inner state shall be copied in upon calling it and any change shall be copied back out onto the PSS entity upon return.

D.4.5 Enumeration types

PSS enumeration types are mapped to C++ unscoped enumeration types (as opposed to enum classes), with the corresponding base type, if any, and with the same set of enum items in the same order and identical names. When specified, explicit numeric constant values for an enum item correspond to the same value in the C++ definition.

For example, the PSS definition:

```
enum color_e {red = 0x10, green = 0x20, blue = 0x30};
```

is mapped to the C++ type as defined by this very same code.

Consequently, enum item names within types used in PSS-to-C++ type binding shall be unique.

D.4.6 Export-function context handle

In C++, a void pointer is used as the type for export-function context handles.

Example D.4 shows the definition of a PSS function that is exported. Example D.5 shows the C prototype for that function.

```
function void exp_func(int val) {
}
export target function exp_func;
```

Example D.4—PSS-defined exported function

Copyright © 2026 Accellera. All rights reserved.

```
void exp_func(void *context, int val);
```

Example D.5—C prototype for exported function

D.5 SystemVerilog language bindings

D.5.1 Function names

PSS implementations shall support mapping a PSS function name to an identical function or task name in SystemVerilog, ignoring the hierarchical path in PSS. PSS implementations may define additional mapping schemes for function names.

D.5.2 Primitive types

The mapping between the PSS primitive types and SystemVerilog types for both parameter and return types is specified in Table D7.

Table D.7—Mapping PSS primitive types and SystemVerilog types

PSS type	SystemVerilog type
string	**string**
bool	**bit**
chandle	**chandle**
bit (1-8-bit domain)	**byte unsigned**
bit (9-16-bit domain)	**shortint unsigned**
bit (17-32-bit domain)	**int unsigned**
bit (33-64-bit domain)	**longint unsigned**
int (1-8-bit domain)	**byte**
int (9-16-bit domain)	**shortint**
int (17-32-bit domain)	**int**
int (33-64-bit domain)	**longint**
float32	**shortreal**
float64	**real**

PSS functions designated with the **target** qualifier (see 20.4.1) may be mapped either to tasks or functions in SystemVerilog, and shall be mapped to tasks by default. PSS **solve** functions shall be mapped to SystemVerilog functions. If neither platform qualifier is used, the default mapping shall be to a function. PSS functions that are mapped to SystemVerilog tasks may not be called on the solve platform.

When a PSS function is mapped to a SystemVerilog function, the return type (if any) and arguments of the SystemVerilog function shall correspond to those of the PSS function prototype.

Copyright © 2026 Accellera. All rights reserved.

When a PSS function is mapped to a SystemVerilog task, the following apply:

a) If the PSS function is a **void** function, then all arguments of the SystemVerilog task shall correspond to the PSS prototype:

$$f(D_0\ T_0\ p_0,\ D_1\ T_1\ p_1,\ \ldots,\ D_n\ T_n\ p_n);\ \Rightarrow\ t(D_0\ T'_0\ p_0,\ D_1\ T'_1\ p_1,\ \ldots,\ D_n\ T'_n\ p_n);$$

b) If the PSS function returns a value, then the first argument of the SystemVerilog task shall be an output of the type corresponding to the return value. All other arguments shall correspond accordingly:

$$T_r\ f(D_0\ T_0\ p_0,\ D_1\ T_1\ p_1,\ \ldots,\ D_n\ T_n\ p_n);\ \Rightarrow\ t(\text{output}\ T'_r\ p_r,\ D_0\ T'_0\ p_0,\ D_1\ T'_1\ p_1,\ \ldots,\ D_n\ T'_n\ p_n);$$

D.5.3 Numeric value mapping

When a numeric type or value is passed from PSS to SystemVerilog, the value shall be expanded or truncated according to SystemVerilog rules (IEEE 1800-2017, section 10.7), treating the SystemVerilog type as the left-hand side of an assignment statement where the PSS value is the right-hand side.

When a numeric type of value is passed from SystemVerilog to PSS, the value shall be expanded or truncated according to the rules in 8.7, treating the SystemVerilog type as the right-hand side of an assignment statement where the PSS value is the left-hand side.

D.5.4 Arrays

Fixed-size arrays in PSS are mapped to SystemVerilog dynamic arrays of corresponding type. Arrays are passed by value between PSS and SystemVerilog.

D.5.5 Lists

Lists in PSS are mapped to SystemVerilog queues of the corresponding type. As with arrays (D.5.4), lists are passed by value between PSS and SystemVerilog. The list may contain any of the primitive types (D.5.2) as well as structs (D.5.6) and enumeration types (D.5.7).

D.5.6 Structs

PSS **struct** types are mapped to classes in SystemVerilog with fields whose types correspond and whose names match. Values of all fields are deep-copied between mapped elements.

The following also apply:

a) The target SystemVerilog class shall contain all fields present in the PSS **struct**. The target SystemVerilog class may be derived from a base class type.

b) Inheritance relationships may or may not be the same across the boundary. Whether the PSS **struct** is derived from a base type has no bearing on whether the SystemVerilog class to which it is mapped is derived from a similar (or any) type.

c) Passing inheritance hierarchies with shadowed fields is not supported.

d) Tools shall ignore the containing namespace of mapped structs.

Copyright © 2026 Accellera. All rights reserved.

D.5.7 Enumeration types

A PSS enumeration type is mapped to a SystemVerilog enum type. The integer values of the *enum_items* shall match, but it is not required that the names of the *enum_items* match.

If a PSS enumeration type is passed to or from SystemVerilog, the **enum** value is passed as its integer equivalent, according to D.5.3.

D.5.8 Export-function context handle

In SystemVerilog, a **chandle** is used as the type for export-function context handles.

Example D.6 shows the definition of a PSS function that is exported. Example D.7 shows the SystemVerilog prototype for that function.

```
function void exp_func(int val) {
}
export target function exp_func;
```

Example D.6—PSS-defined exported function

```
function void exp_func(chandle context, int val);
```

Example D.7—SystemVerilog prototype for exported function

Copyright © 2026 Accellera. All rights reserved.

Annex E
(informative)
Solution space

Once a PSS model has been specified, the elements of the model shall be processed in some way to ensure that resulting scenarios accurately reflect the specified behaviors. This annex describes the steps a processing tool may take to analyze a portable stimulus description and create a (set of) scenario(s). See also Clause 14.

a) Identify root action:
   1) Specified by the user.
   2) Unless otherwise specified, the designated root action shall be located in the root component. By default, the root component shall be **`pss_top`**.
   3) If the specified root action is an atomic action, consider it to be the initial action traversed in an implicit **activity** statement.
   4) If the specified root action is a compound action:
      i) Identify all **bind** statements in the activity and bind the associated object(s) accordingly. Identify all resulting scheduling dependencies between bound actions.
      ii) For every compound action traversed in the activity, expand its activity to include each sub-action traversal in the overall activity to be analyzed.
      iii) Identify scheduling dependencies among all action traversals declared in the activity and add to the scheduling dependency list identified in a.4.i.

b) For each action traversed in the activity:
   1) For each resource locked or shared (i.e., claimed) by the action:
      i) Identify the resource pool of the appropriate type to which the resource reference may be bound.
      ii) Identify all other action(s) claiming a resource of the same type that is bound to the same pool.
      iii) Each resource object instance in the resource pool has an built-in **`instance_id`** field that is unique for that pool.
      iv) The algebraic constraints for evaluating field(s) of the resource object are the union of the constraints defined in the resource object type and the constraints in all actions ultimately connected to the resource object.
      v) Identify scheduling dependencies enforced by the claimed resource and add these to the set of dependencies identified in a.4.i.
         1. If an action locks a resource instance, no other action claiming that same resource instance may be scheduled concurrent with the locking action.
         2. If actions scheduled concurrently collectively attempt to lock more resource instances than are available in the pool, an error shall be generated.
         3. If the resource instance is not locked, there are no scheduling implications of sharing a resource instance.
   2) For each flow object declared in the action that is not already bound:
      i) If the flow object is not explicitly bound to a corresponding flow object, identify the object pool(s) of the appropriate type to which the flow object may be bound.
      ii) The algebraic constraints for evaluating field(s) of the flow object are the union of the constraints defined in flow object type and the constraints in all actions ultimately connected to the flow object.
      iii) Identify all other explicitly-traversed actions bound to the same pool that:

Copyright © 2026 Accellera. All rights reserved.

1. Declare a matching object type with consistent data constraints,
2. Meet the scheduling constraints from b.1.v, and
3. Are scheduled consistent with the scheduling constraints implied by the type of the flow object.

iv) The set of explicitly-traversed actions from b.2.iii shall compose the *inferencing candidate list* (*ICL*).

v) If no explicitly traversed action appears in the ICL, then an anonymous instance of each action type bound to the pool from b.2.i shall be added to the ICL.

vi) If the ICL is empty, an error shall be generated.

vii) For each element in the ICL, perform step b.2 until no actions in the ICL have any unbound flow object references or the tool's inferencing limit is reached (see c).

c) If the tool reaches the maximum inferencing depth, it shall infer a terminating action if one is available. Given the set of actions, flow and resource objects, scheduling and data constraints, and associated ICLs, pick an instance from the ICL and a value for each data field in the flow object that satisfies the constraints and bind the flow object reference from the action to the corresponding instance from the ICL.

Copyright © 2026 Accellera. All rights reserved.

Annex F
(normative)
Formal semantics of behavioral coverage

F.1 General

This annex describes the formal semantics of PSS behavioral coverage. PSS data coverage is not discussed here.

The semantics description is based on the abstract syntax, which includes only basic constructs and ignores the metalanguage features, such as monitor extension, overriding, and inheritance.

It is assumed that all action and monitor handles are unique and defined at the top level. The handles become unique if their scope and constant indexing (if any) are considered part of their name.

The semantics of expressions, **foreach** and **forall** constraints is assumed to be known.

covergroup semantics is completely defined by a mapping of action handles into action executions at the first match point of the top-level monitor attempts of a cover statement; this is why there is no need to describe the **covergroup** construct separately.

F.2 Definitions and notation

Throughout this annex, the notation described in this subclause will be used.

- A *cover statement* is denoted by C. The top-level scenario of a cover statement is denoted as *scenario(C)*.
- A *scenario* is denoted by $s,s_1,s_2,\ldots$.
- The *set of all action executions* is denoted by $\mathcal{X}$.
- Individual *action executions* are denoted by $x,x_1,x_2,\ldots$.
- *Sets of action executions* are denoted by $X,X_1,X_2,\ldots$.
- An *action handle* is denoted as $h,h_0,h_1,\ldots$ An action handle may be explicitly defined (*action_type h;*) or anonymous (`do` *action_type*).
- *Sets of action handles* are denoted by $H,H_1,H_2,\ldots$.
- *Monitor handle* is denoted by m (also includes an anonymous handle `do` M, where M is monitor type). *scenario(m)* denotes the monitor scenario.
- *Boolean predicates* on action executions are denoted by $p,p_0,p_1,\ldots$.
- A *scenario realization* is denoted by $r,r_1,r_2,\ldots$.
- *Sets of scenario realizations* are denoted by $R,R_1,R_2,\ldots$.
- A *time instant* (non-negative integer number) is denoted by $t,t_0,t_1,\ldots$.
- A *start time* of an action execution x or of a scenario realization r is denoted by *begin(x)* and *begin(r)*, respectively.
- An *end time* of an action execution x or of a scenario realization r is denoted by *end(x)* and *end(r)*, respectively.
- A *domain* of function f is denoted by dom(f).
- An *image (codomain)* of function f is denoted by im(f).

Copyright © 2026 Accellera. All rights reserved.

— A *scenario realization function* is denoted by *realizations*(*s,t,R*), where *s* is a scenario, *t* is a checkpoint and *R* is an input set of action realizations (see F.4.2).

F.3 Abstract syntax

The formal semantics is based on an abstract syntax of behavioral coverage rather than on the full PSS BNF. The abstract syntax facilitates separation of derived monitor activity statements from the basic ones. The scenario realization is defined explicitly only for basic statements.

F.3.1 Abstract grammar

The abstract grammar is based on the following symbols considered as terminals:

— *h* is an action handle, either explicit or anonymous.

— *p* is a Boolean expression corresponding to an atomic algebraic constraint.

Every action handle in an action traversal has an inline constraint. Inline constraint expression *true* is equivalent to the omitted inline constraint.

Scenario:

{}

| *h with p*

| *concat* { *s; s;* }

| *eventually s;*

| *select* {*s; s;* }

| *schedule* {*s;s;*}

| *overlap* {*scenario_set* }

| *m with p*

| *s constraint* { *p* }

scenario_set::=s; s; | *scenario_set; s*

{} above denotes an empty sequence.

F.3.2 Derived forms

Derived forms are considered shortcuts. Their rewriting below is defined using ≡ notation.

F.3.2.1 Constraints

The constraints are Boolean predicates on action executions (see F.4.1).

Copyright © 2026 Accellera. All rights reserved.

— Implication constraint:
$p_1 \rightarrow p_2 \equiv !p_1 \| p_2$
— If-else constraints:
`if` (p_1) $p_2 \equiv !p_1 \| p_2$
`if` (p_1) p_2 `else` $p_3 \equiv p_1 \&\& p_2 \| !p_1 \&\& p_3$

A set of internal constraints $\{p_1,p_2,...\}$ is interpreted as $p_1 \&\& p_2 \&\& ....$

A set of standalone constraints `constraint` $\{p_1,p_2,...\}$ is interpreted as `constraint` p_1; `constraint` p_2; ....

F.3.2.2 Scenarios

`concat { `s`; } ≡ `s
`concat {`s_1`; ...,`s_{n-1}`; `s_n`;} = concat { concat { `s_1`; ...,`s_{n-1}`}; `s_n`;}, `$n \geq 3$
`select {`s_1`; ...,`s_{n-1}`; `s_n`;} = select { select { `s_1`; ...,`s_{n-1}`}; `s_n`;}, `$n \geq 3$
`schedule { `s`; } ≡ `s
`schedule {`s_1`; ...,`s_{n-1}`; `s_n`;} = schedule { schedule { `s_1`; ...,`s_{n-1}`}; `s_n`;}, `$n \geq 3$
`overlap { `s`; } ≡ `s
`overlap {{}, `s_2`,...} ≡ overlap {`s_2`,...}`
`overlap { ...,`s`, {}} ≡ overlap {...,`s`}`
`overlap {...,`s_1`, {}, `s_2`,...} ≡ overlap {...,`s_1`,`s_2`,...}`
`sequence { `s`; {} } ≡ `s
`sequence { `s_1`;`s_2`; } ≡ concat { `s_1`; eventually `s_2`;}`
`sequence {`s_1`; ...,`s_{n-1}`; `s_n`;} ≡ sequence { sequence { `s_1`; ...,`s_{n-1}`}; `s_n`;}, `$n \geq 3$

F.4 Semantics

F.4.1 Action execution model

All action executions form a set $\mathcal{X}$. Attribute f of type T (e.g., `bool`, `int`) of an action of type a may be considered as a function of a signature $f{:}a \rightarrow T$, i.e., as a function receiving an argument of type a, and returning a value of type T. In the following example:

```
action write { rand int core; }
```

action type `write` defines attribute `core` with the integer domain. Here f= `core`, a= `write`, and T= `int`. The function `core` is interpreted on specific executions of action `write`, and may assume values 0, 1, ...

Action type a defines a set of this action attributes $f_1{:}a \rightarrow T_1$, ..., $f_k{:}a \rightarrow T_k$, k=1, 2, ...

An algebraic constraint may be considered as a predicate $p{:}a_1 \rightarrow a_2 \rightarrow ... \rightarrow a_n \rightarrow$ `bool`, i.e., as a Boolean function whose arguments are action types a_1, a_2, …, a_n, n=1, 2, ... More precisely, a predicate is a function of the following form:

$$p(h_1, h_2, \dots, h_n) = p\left(f_{11}(h_1), \dots, f_{1k_1}(h_1), f_{21}(h_2), \dots, f_{2k_2}(h_2), \dots, f_{n1}(h_n), \dots, f\left(h_{nk_n}\right)\right), k_1, k_2, \dots, k_n = 1, 2, \dots$$

In this case, h_1,...,h_n are action handles (explicit or anonymous) of the following types:
$a_1, \dots, a_n$, $i = 1, \dots, n$, and $f_{i1}, \dots, f_{ik_i}$

Copyright © 2026 Accellera. All rights reserved.

are their attributes.

In the following example:

```
action read { rand int core; rand bit[32] addr; rand bool locked; }
monitor m { write w; read r;
    activity { w; r with core == 1 && locked && addr == w.addr; }
}
```

the predicate `core == 1 && locked && addr == w.addr` has the form $p(f_{11}(h_1), f_{21}(h_2), f_{22}(h_2), f_{23}(h_2))$, where h_1= `write`, h_2= `read`, $f_{11}(h_1)$= `w.addr`, $f_{21}(h_2)$= `r.core`, $f_{22}(h_2)$ = `r.locked`, $f_{23}(h_2)$= `r.addr`, and $p(x,y,z,t)$=(y==1 && z && t==x). Here x,y,z,t are arguments of predicate p.

An action object x has its type a, and thus all attributes defined by a, and also its beginning and end times, denoted by $begin(x)$ and $end(x)$, correspondingly. The beginning and end times are non-negative integers, $begin(x) < end(x)$. Intuitively, action execution x spans from $begin(x)$ until and not including $end(x)$.

F.4.2 Scenario realization

A *scenario realization* is a one-to-one function r (map) from a set of action handles $H=\{h_1,...,h_n\}$ into a set of action executions $X=\{x_1,...,x_n\} \subset \mathcal{X}, n > 0$: $\{h_1,...,h_n\} \rightarrow \{x_1,...,x_n\}$; X is a *scenario realization domain*, H is a *scenario realization codomain* or *image*. For scenario realizations the following notation will be used: $r = \{x_1 \mapsto h_1,...,x_n \mapsto h_n\}$.

As an example, consider a scenario defined by monitor `m` (F.4.1) and the action execution trace shown in Figure 1.

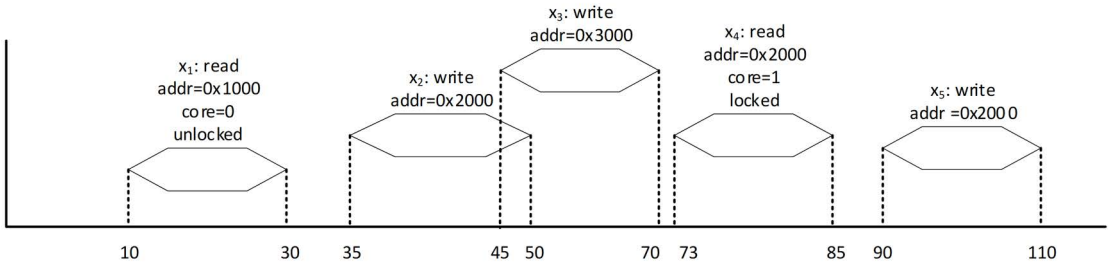


Figure F.1—Action execution trace

Here, the set of all action execution $\mathcal{X}=\{x_1,x_2,x_3,x_4,x_5\}$. One scenario realization r is {`do write` ↦ x_2, `r` ↦ x_4}. The scenario realization domain dom (r)= H= {`do write`,`r`} contains one anonymous (`do write`) and one explicit handle (`r`); the scenario realization codomain contains two action executions: im (r)=X=$\{x_2,x_4\}$.

By the *disjoint union* of two sets $A \sqcup B$ is understood their union $A \cup B$ provided these sets do not intersect: $A \cap B = \emptyset$. The notion of the disjoint union may be extended into two functions f and g, denoted as $f \sqcup g$, provided their domains are disjoint ($dom(f) \cap dom(g) = \emptyset$) as follows: $dom(f \sqcup g) = dom(f) \sqcup dom(g)$ and $f \sqcup g(x) = f(x)$ if $x \in dom(f)$ and $f \sqcup g(x) = g(x)$ if $x \in dom(g)$. The disjoint union of scenario realizations $r_1 \sqcup r_2$ is a special case of the disjoint union of functions.

An *application of a predicate* p to a scenario realization r, denoted as $p[r]$ is defined by replacing handle attributes with their values in the corresponding action executions. Thus for the constraint

Copyright © 2026 Accellera. All rights reserved.

p `= (core == 1) (&& locked && addr == w.addr)` in the monitor m above and the above realization r, $p[r]$ `= (1 == 1) && locked && (0x2000 == 0x2000)` (which evaluates to true).

The *start time of a scenario realization r* is the minimal begin time of all its action executions:

$$begin(r) = \min_{x \in \mathrm{im}\, r} begin(x)$$

The *end (or match) time* of a scenario realization *r* is the maximal time of all its action executions:

$$end(r) = \max_{x \in \mathrm{im}\, r} end(x)$$

F.4.3 Coverage semantics

The set of action objects is *covered* by cover statement *C* iff there exists t=0,1,... such that there is a successful top-level attempt starting at time *t*. The latter means $R=\{r \mid r \in realizations(scenario(C),t,\emptyset)$ and $t=begin(r)\} \neq \emptyset$. In other words, the attempt scenario has at least one realization starting at time *t*.

The *scenario realization function realizations*(*s*,*t*, *R*) defines the set of realizations of scenario *s* with the checkpoint *t* with the input set of realizations *R*. It is defined recursively as follows (with the convention that a union of zero number of sets is ∅):

$$realizations(\{\}) = \{\emptyset\}$$

i.e., the only realization of an empty sequence is empty (does not contain any action executions).

$$realizations(h\ with\ p, t, R) = \cup_{r \in R} \bigcup_{x} \{(h \mapsto x) | p[r \sqcup \{h \mapsto x\}]\}$$

where $r \in R$, $x \in X$, $begin(x) \geq t$ and there is no $y \in X$, $t \leq begin(y) < begin(x)$. This means that the realization of an action traversal scenario corresponds to all action executions of an appropriate type and appropriately constrained and closest to the checkpoint *t*. It is a syntax error if $\mathrm{supp}(p) \not\subset \mathrm{dom}(r) \cup \{h\}$. Here $\mathrm{supp}(p)$ denotes the set of variables on which the constraint *p* depends.

$$realizations(concat\{s_1, s_2\}, t, R) = \cup_{r_1} \bigcup_{r_2} \{r_1 \sqcup r_2\}$$

where $r_1 \in realizations(s_1, t, R)$, $r_2 \in realizations(s_2, end(r_1), \{r_1\}))$ if $r_1 \neq \emptyset$, and $r_2 \in realizations(s_2, t, R)$, otherwise. This means that every realization of *concat* is a realization of its first argument combined with a realization of its second argument with the checkpoint at the end time of the realization of the first argument.

$$realizations(eventually\ s, t, R) = \bigcup_{\tau \geq t} \{realizations(s, \tau, R)\}$$

This means that the realizations of scenario *s* are collected at each time instant starting from *t*.

$$realizations(select\{s_1, s_2\}, t, R) = realizations(s_1, t, R) \cup realizations(s_2, t, R)$$

This means that the realizations of *select* are realizations of one of its scenarios.

$$realizations(schedule\{s_1, s_2\}, t, R) = \bigcup_{r_1, r_2} \{r_1 \sqcup r_2\}$$

where $r_1 \in realizations(eventually\ s_1, t, R)$, $r_2 \in realizations(eventually\ s_2, t, R)$ and $\mathrm{dom}(r_1) \cap \mathrm{dom}(r_2) = \emptyset$. This means that the realizations of *schedule* are combinations of realizations of its arguments with the checkpoints at *t* or in the future; the combined realizations do not have common action executions.

Copyright © 2026 Accellera. All rights reserved.

$$realizations(overlap\{s_1, \dots, s_n\}, t, R) = \bigcup_{r_1,\dots,r_n} \{r_1 \sqcup \dots \sqcup r_n\}$$

where $r_1 \in realizations(s_1,t,R),\dots,r_n \in realizations(s_n,t,R)$, $n \geq 2$ and for all $1 \leq i < j \leq n$ $\mathrm{dom}(r_i) \cap \mathrm{dom}(r_j) = \emptyset$ and $\max(begin(r_1),\dots,begin(r_n)) < \min(end(r_1),\dots,end(r_n))$. Here it is assumed for empty realizations that $begin(\emptyset) = -\infty$ and $end(\emptyset) = +\infty$. This definition is similar to the definition of *schedule*, but in addition, it requires that the member realization windows overlap.

$$realizations(m, t, R) = realizations(scenario(m), t, R)$$

i.e., the realizations of a monitor are the realizations of its top-level scenario.

$$realizations(s\ constraint\ \{p\}, t, R\} = \bigcup_{r \in realizations(s,t,R)} \{r\}$$

where either $\mathrm{supp}(p) \not\subset \mathrm{dom}(r)$ or $p[r]$ evaluates to true. Here $\mathrm{supp}(p)$ denotes the set of handles mentioned in p. This means that the realization should satisfy the imposed constraint, maybe vacuously.

Note that the scenario realization function is defined only when the beginning and the end times of its appropriate member realizations are defined.

As an illustration, consider the set of the realizations of the top-level scenario of monitor cover statement

```
cov: cover { m}
```

where `m` is defined in F.4.1 as

```
monitor m { write w; read r;
    activity { w; r with core == 1 && locked && addr == w.addr; }
}
```

for the action execution trace shown in Figure 2 (reproduction of Figure 1) for the checkpoint t=35.

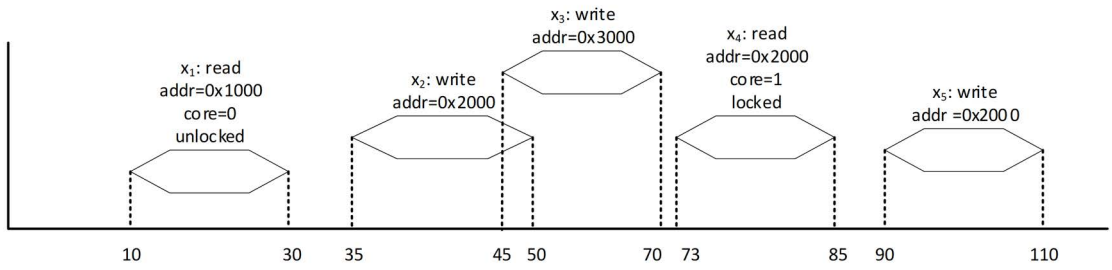


Figure F.2—Action execution trace

To be brief, we will write *a* instead of

```
r with core == 1 && locked && addr == w.addr
```

realizations(*m*,35,∅) = *realizations*(`sequence { w; r with core == 1 && locked && addr == w.addr; }`,35,∅) = (abbreviation)

realizations(`sequence` {*w*;*a* }) = (definition of **sequence**)

realizations(`concat { w;eventually` *a*;},35,∅).

Copyright © 2026 Accellera. All rights reserved.

According to the definition of **concat**, its scenario realization function is computed as:

$$\bigcup_{r_1} \bigcup_{r_2} \{r_1 \sqcup r_2\}$$

where $r_1 \in$ *realizations*$(w,t,\emptyset)$,$r_2 \in$ *realizations*(`eventually` a,*end*(r_1),$\{r_1\}$)).

realizations$(w,t,\emptyset) = \{w \mapsto x_2\}$. Indeed, action executions of type write are x_2, x_3, and x_5, 35=*begin*(x_2) < *begin*(x_3) < *begin*(x_5) and the inline constraint is void. Thus, *realizations*$(w,t,\emptyset)$ consists of the only realization $r_1 = w \mapsto x_2$ and its endpoint *begin*$(r_1) =$ *begin*$(x_2) = 50$.

According to the definition:

$$realizations(\texttt{eventually}\ a, 50, \texttt{w} \mapsto x_2) = \bigcup_{\tau \geq 50} realizations(a, \tau, \texttt{w} \mapsto x_2).$$

It is easy to see that *realizations*$(a, 50, w \mapsto x_2) = \{r \mapsto x_4\}$. Indeed, x_4 is the only `read` action execution starting at or after time 50 and `r.core == 1 && r.locked && r.addr == w.addr` evaluates to true when r maps into x_4 and w maps into x_2. It is obvious that *realizations*$(a,50,w \mapsto x_2) = ... =$ *realizations*$(a,73,w \mapsto x_2) =$ *realizations*$(a,50,w \mapsto x_2) = \{r \mapsto x_4\}$ and *realizations*$(a,74,w \mapsto x_2) = p(a,75,w \mapsto x_2) = ... = \emptyset$ (no more read actions after time 73). Thus, *realizations*(*eventually* a,50,$w \mapsto x_2$) $= \{r \mapsto x_4\} = \{r_2\}$. The final result is $\{\texttt{w} \mapsto x_2, \texttt{r} \mapsto x_4\}$.

Copyright © 2026 Accellera. All rights reserved.