



TOSCA Simple Profile in YAML Version 1.1

OASIS Standard

30 January 2018

Specification URIs

This version:

<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.1/os/TOSCA-Simple-Profile-YAML-v1.1-os.pdf> (Authoritative)
<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.1/os/TOSCA-Simple-Profile-YAML-v1.1-os.html>
<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.1/os/TOSCA-Simple-Profile-YAML-v1.1-os.docx>

Previous version:

<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.1/csprd01/TOSCA-Simple-Profile-YAML-v1.1-csprd01.pdf> (Authoritative)
<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.1/csprd01/TOSCA-Simple-Profile-YAML-v1.1-csprd01.html>
<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.1/csprd01/TOSCA-Simple-Profile-YAML-v1.1-csprd01.docx>

Latest version:

<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.1/TOSCA-Simple-Profile-YAML-v1.1.pdf> (Authoritative)
<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.1/TOSCA-Simple-Profile-YAML-v1.1.html>
<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.1/TOSCA-Simple-Profile-YAML-v1.1.docx>

Technical Committee:

OASIS Topology and Orchestration Specification for Cloud Applications (TOSCA) TC

Chairs:

Paul Lipton (paul.lipton@ca.com), CA Technologies
John Crandall (jcrandal@brocade.com), Brocade Communications Systems

Editors:

Matt Rutkowski (mrutkows@us.ibm.com), IBM
Luc Boutier (luc.boutier@fastconnect.fr), FastConnect

Related work:

This specification is related to:

- *Topology and Orchestration Specification for Cloud Applications Version 1.0*. Edited by Derek Palma and Thomas Spatzier. 25 November 2013. OASIS Standard. <http://docs.oasis-open.org/tosca/TOSCA/v1.0/os/TOSCA-v1.0-os.html>.

Declared XML namespace:

- <http://docs.oasis-open.org/tosca/ns/simple/yaml/1.1>

Abstract:

This document defines a simplified profile of the TOSCA Version 1.0 specification in a YAML rendering which is intended to simplify the authoring of TOSCA service templates. This profile

defines a less verbose and more human-readable YAML rendering, reduced level of indirection between different modeling artifacts as well as the assumption of a base type system.

Status:

This document was last revised or approved by the membership of OASIS on the above date. The level of approval is also listed above. Check the “Latest version” location noted above for possible later revisions of this document. Any other numbered Versions and other technical work produced by the Technical Committee (TC) are listed at https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=tosca#technical.

TC members should send comments on this specification to the TC’s email list. Others should send comments to the TC’s public comment list, after subscribing to it by following the instructions at the “Send A Comment” button on the TC’s web page at <https://www.oasis-open.org/committees/tosca/>.

This OASIS Standard is provided under the [RF on Limited Terms](#) Mode of the [OASIS IPR Policy](#), the mode chosen when the Technical Committee was established. For information on whether any patents have been disclosed that may be essential to implementing this specification, and any offers of patent licensing terms, please refer to the Intellectual Property Rights section of the TC’s web page (<https://www.oasis-open.org/committees/tosca/ipr.php>).

Note that any machine-readable content ([Computer Language Definitions](#)) declared Normative for this Work Product is provided in separate plain text files. In the event of a discrepancy between any such plain text file and display content in the Work Product's prose narrative document(s), the content in the separate plain text file prevails.

Citation format:

When referencing this specification the following citation format should be used:

[TOSCA-Simple-Profile-YAML-v1.1]

TOSCA Simple Profile in YAML Version 1.1. Edited by Matt Rutkowski and Luc Boutier. 30 January 2018. OASIS Standard. <http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.1/os/TOSCA-Simple-Profile-YAML-v1.1-os.html>. Latest version: <http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.1/TOSCA-Simple-Profile-YAML-v1.1.html>.

Notices

Copyright © OASIS Open 2018. All Rights Reserved.

All capitalized terms in the following text have the meanings assigned to them in the OASIS Intellectual Property Rights Policy (the "OASIS IPR Policy"). The full [Policy](#) may be found at the OASIS website.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published, and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this section are included on all such copies and derivative works. However, this document itself may not be modified in any way, including by removing the copyright notice or references to OASIS, except as needed for the purpose of developing any document or deliverable produced by an OASIS Technical Committee (in which case the rules applicable to copyrights, as set forth in the OASIS IPR Policy, must be followed) or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by OASIS or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and OASIS DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY OWNERSHIP RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

OASIS requests that any OASIS Party or any other party that believes it has patent claims that would necessarily be infringed by implementations of this OASIS Committee Specification or OASIS Standard, to notify OASIS TC Administrator and provide an indication of its willingness to grant patent licenses to such patent claims in a manner consistent with the IPR Mode of the OASIS Technical Committee that produced this specification.

OASIS invites any party to contact the OASIS TC Administrator if it is aware of a claim of ownership of any patent claims that would necessarily be infringed by implementations of this specification by a patent holder that is not willing to provide a license to such patent claims in a manner consistent with the IPR Mode of the OASIS Technical Committee that produced this specification. OASIS may include such claims on its website, but disclaims any obligation to do so.

OASIS takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on OASIS' procedures with respect to rights in any document or deliverable produced by an OASIS Technical Committee can be found on the OASIS website. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this OASIS Committee Specification or OASIS Standard, can be obtained from the OASIS TC Administrator. OASIS makes no representation that any information or list of intellectual property rights will at any time be complete, or that any claims in such list are, in fact, Essential Claims.

The name "OASIS" is a trademark of [OASIS](#), the owner and developer of this specification, and should be used only to refer to the organization and its official outputs. OASIS welcomes reference to, and implementation and use of, specifications, while reserving the right to enforce its marks against misleading uses. Please see <https://www.oasis-open.org/policies-guidelines/trademark> for above guidance.

Table of Contents

Table of Examples	7
Table of Figures	7
1 Introduction	8
1.0 IPR Policy	8
1.1 Objective	8
1.2 Summary of key TOSCA concepts	8
1.3 Implementations	9
1.4 Terminology	9
1.5 Notational Conventions	9
1.6 Normative References	9
1.7 Non-Normative References	10
1.8 Glossary	10
2 TOSCA by example	12
2.1 A “hello world” template for TOSCA Simple Profile in YAML	12
2.2 TOSCA template for a simple software installation	14
2.3 Overriding behavior of predefined node types	16
2.4 TOSCA template for database content deployment	16
2.5 TOSCA template for a two-tier application	18
2.6 Using a custom script to establish a relationship in a template	20
2.7 Using custom relationship types in a TOSCA template	22
2.8 Defining generic dependencies between nodes in a template	23
2.9 Describing abstract requirements for nodes and capabilities in a TOSCA template	24
2.10 Using node template substitution for model composition	28
2.11 Using node template substitution for chaining subsystems	32
2.12 Grouping node templates	37
2.13 Using YAML Macros to simplify templates	39
2.14 Passing information as inputs to Nodes and Relationships	40
2.15 Topology Template Model versus Instance Model	42
2.16 Using attributes implicitly reflected from properties	42
3 TOSCA Simple Profile definitions in YAML	44
3.1 TOSCA Namespace URI and alias	44
3.2 Parameter and property types	45
3.3 Normative values	54
3.4 TOSCA Metamodel	55
3.5 Reusable modeling definitions	55
3.6 Type-specific definitions	82
3.7 Template-specific definitions	98
3.8 Topology Template definition	110
3.9 Service Template definition	115
4 TOSCA functions	127
4.1 Reserved Function Keywords	127
4.2 Environment Variable Conventions	127
4.3 Intrinsic functions	130

4.4	Property functions	131
4.5	Attribute functions	133
4.6	Operation functions.....	134
4.7	Navigation functions	135
4.8	Artifact functions	136
4.9	Context-based Entity names (global)	138
5	TOSCA normative type definitions	139
5.1	Assumptions	139
5.2	TOSCA normative type names.....	139
5.3	Data Types.....	140
5.4	Artifact Types.....	147
5.5	Capabilities Types	150
5.6	Requirement Types	159
5.7	Relationship Types	159
5.8	Interface Types	162
5.9	Node Types.....	167
5.10	Group Types	178
5.11	Policy Types	179
6	TOSCA Cloud Service Archive (CSAR) format.....	181
6.1	Overall Structure of a CSAR.....	181
6.2	TOSCA Meta File.....	181
6.3	Archive without TOSCA-Metadata.....	182
7	TOSCA workflows	183
7.1	Normative workflows.....	183
7.2	Declarative workflows	183
7.3	Imperative workflows	187
7.4	Making declarative more flexible and imperative more generic	199
8	TOSCA networking.....	202
8.1	Networking and Service Template Portability.....	202
8.2	Connectivity Semantics	202
8.3	Expressing connectivity semantics.....	203
8.4	Network provisioning	205
8.5	Network Types.....	209
8.6	Network modeling approaches	214
9	Non-normative type definitions.....	219
9.1	Artifact Types.....	219
9.2	Capability Types	219
9.3	Node Types.....	221
10	Component Modeling Use Cases.....	224
11	Application Modeling Use Cases.....	231
11.1	Use cases	231
12	TOSCA Policies.....	271
12.1	A declarative approach.....	271
12.2	Consideration of Event, Condition and Action	271
12.3	Types of policies	271

12.4 Policy relationship considerations	272
12.5 Use Cases	273
13 Conformance	276
13.1 Conformance Targets	276
13.2 Conformance Clause 1: TOSCA YAML service template	276
13.3 Conformance Clause 2: TOSCA processor.....	276
13.4 Conformance Clause 3: TOSCA orchestrator	276
13.5 Conformance Clause 4: TOSCA generator	277
13.6 Conformance Clause 5: TOSCA archive	277
Appendix A. Known Extensions to TOSCA v1.0.....	278
A.1 Model Changes	278
A.2 Normative Types	278
Appendix B. Acknowledgments	280
Appendix C. Revision History.....	281

Table of Examples

Example 1 - TOSCA Simple "Hello World"	12
Example 2 - Template with input and output parameter sections	13
Example 3 - Simple (MySQL) software installation on a TOSCA Compute node	14
Example 4 - Node Template overriding its Node Type's "configure" interface	16
Example 5 - Template for deploying database content on-top of MySQL DBMS middleware	17
Example 6 - Basic two-tier application (web application and database server tiers)	18
Example 7 - Providing a custom relationship script to establish a connection	21
Example 8 - A web application Node Template requiring a custom database connection type	22
Example 9 - Defining a custom relationship type	23
Example 10 - Simple dependency relationship between two nodes	23
Example 11 - An abstract "host" requirement using a node filter	24
Example 12 - An abstract Compute node template with a node filter	25
Example 13 - An abstract database requirement using a node filter	26
Example 14 - An abstract database node template	27
Example 15 - Referencing an abstract database node template	29
Example 16 - Using substitution mappings to export a database implementation	31
Example 17 - Declaring a transaction subsystem as a chain of substitutable node templates	33
Example 18 - Defining a TransactionSubsystem node type	34
Example 19 - Implementation of a TransactionSubsystem node type using substitution mappings	36
Example 20 - Grouping Node Templates for possible policy application	37
Example 21 - Grouping nodes for anti-colocation policy application	38
Example 22 - Using YAML anchors in TOSCA templates	40
Example 23 - Properties reflected as attributes	42

Table of Figures

Figure 1: Using template substitution to implement a database tier	29
Figure 2: Substitution mappings	31
Figure 3: Chaining of subsystems in a service template	33
Figure 4: Defining subsystem details in a service template	35
Figure-5: Typical 3-Tier Network	206
Figure-6: Generic Service Template	215
Figure-7: Service template with network template A	215
Figure-8: Service template with network template B	216

1 Introduction

1.0 IPR Policy

This OASIS Standard has been developed under the [RF on Limited Terms](#) Mode of the [OASIS IPR Policy](#), the mode chosen when the Technical Committee was established.

For information on whether any patents have been disclosed that may be essential to implementing this specification, and any offers of patent licensing terms, please refer to the Intellectual Property Rights section of the TC's web page (<https://www.oasis-open.org/committees/tosca/ipr.php>).

1.1 Objective

The TOSCA Simple Profile in YAML specifies a rendering of TOSCA which aims to provide a more accessible syntax as well as a more concise and incremental expressiveness of the TOSCA DSL in order to minimize the learning curve and speed the adoption of the use of TOSCA to portably describe cloud applications.

This proposal describes a YAML rendering for TOSCA. YAML is a human friendly data serialization standard (<http://yaml.org/>) with a syntax much easier to read and edit than XML. As there are a number of DSLs encoded in YAML, a YAML encoding of the TOSCA DSL makes TOSCA more accessible by these communities.

This proposal prescribes an isomorphic rendering in YAML of a subset of the TOSCA v1.0 XML specification ensuring that TOSCA semantics are preserved and can be transformed from XML to YAML or from YAML to XML. Additionally, in order to streamline the expression of TOSCA semantics, the YAML rendering is sought to be more concise and compact through the use of the YAML syntax.

1.2 Summary of key TOSCA concepts

The TOSCA metamodel uses the concept of service templates to describe cloud workloads as a topology template, which is a graph of node templates modeling the components a workload is made up of and as relationship templates modeling the relations between those components. TOSCA further provides a type system of node types to describe the possible building blocks for constructing a service template, as well as relationship type to describe possible kinds of relations. Both node and relationship types may define lifecycle operations to implement the behavior an orchestration engine can invoke when instantiating a service template. For example, a node type for some software product might provide a 'create' operation to handle the creation of an instance of a component at runtime, or a 'start' or 'stop' operation to handle a start or stop event triggered by an orchestration engine. Those lifecycle operations are backed by implementation artifacts such as scripts or Chef recipes that implement the actual behavior.

An orchestration engine processing a TOSCA service template uses the mentioned lifecycle operations to instantiate single components at runtime, and it uses the relationship between components to derive the order of component instantiation. For example, during the instantiation of a two-tier application that includes a web application that depends on a database, an orchestration engine would first invoke the 'create' operation on the database component to install and configure the database, and it would then invoke the 'create' operation of the web application to install and configure the application (which includes configuration of the database connection).

The TOSCA simple profile assumes a number of base types (node types and relationship types) to be supported by each compliant environment such as a 'Compute' node type, a 'Network' node type or a generic 'Database' node type. Furthermore, it is envisioned that a large number of additional types for use in service templates will be defined by a community over time. Therefore, template authors in many cases will not have to define types themselves but can simply start writing service templates that use existing

46 types. In addition, the simple profile will provide means for easily customizing and extending existing
47 types, for example by providing a customized 'create' script for some software.

48 1.3 Implementations

49 Different kinds of processors and artifacts qualify as implementations of the TOSCA simple profile. Those
50 that this specification is explicitly mentioning or referring to fall into the following categories:

- 51 • TOSCA YAML service template (or "service template"): A YAML document artifact containing a
52 (TOSCA) service template (see sections 3.9 "Service template definition") that represents a Cloud
53 application. (see sections 3.8 "Topology template definition")
- 54 • TOSCA processor (or "processor"): An engine or tool that is capable of parsing and interpreting a
55 TOSCA service template for a particular purpose. For example, the purpose could be validation,
56 translation or visual rendering.
- 57 • TOSCA orchestrator (also called orchestration engine): A TOSCA processor that interprets a
58 TOSCA service template or a TOSCA CSAR in order to instantiate and deploy the described
59 application in a Cloud.
- 60 • TOSCA generator: A tool that generates a TOSCA service template. An example of generator is
61 a modeling tool capable of generating or editing a TOSCA service template (often such a tool
62 would also be a TOSCA processor).
- 63 • TOSCA archive (or TOSCA Cloud Service Archive, or "CSAR"): a package artifact that contains a
64 TOSCA service template and other artifacts usable by a TOSCA orchestrator to deploy an
65 application.

66 The above list is not exclusive. The above definitions should be understood as referring to and
67 implementing the TOSCA simple profile as described in this document (abbreviated here as "TOSCA" for
68 simplicity).

69 1.4 Terminology

70 The TOSCA language introduces a YAML grammar for describing service templates by means of
71 Topology Templates and towards enablement of interaction with a TOSCA instance model perhaps by
72 external APIs or plans. The primary currently is on design time aspects, i.e. the description of services to
73 ensure their exchange between Cloud providers, TOSCA Orchestrators and tooling.

74

75 The language provides an extension mechanism that can be used to extend the definitions with additional
76 vendor-specific or domain-specific information.

77 1.5 Notational Conventions

78 The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD
79 NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described
80 in [RFC2119].

81 1.5.1 Notes

- 82 • Sections that are titled "Example" throughout this document are considered non-normative.

83 1.6 Normative References

[RFC2119] S. Bradner, *Key words for use in RFCs to Indicate Requirement Levels*,
<http://www.ietf.org/rfc/rfc2119.txt>, IETF RFC 2119, March 1997.

[TOSCA-1.0] Topology and Orchestration Topology and Orchestration Specification for
Cloud Applications (TOSCA) Version 1.0, an OASIS Standard, 25 November

2013, <http://docs.oasis-open.org/tosca/TOSCA/v1.0/os/TOSCA-v1.0-os.pdf>

[YAML-1.2]	YAML, Version 1.2, 3rd Edition, Patched at 2009-10-01, Oren Ben-Kiki, Clark Evans, Ingy döt Net http://www.yaml.org/spec/1.2/spec.html
[YAML-TS-1.1]	Timestamp Language-Independent Type for YAML Version 1.1, Working Draft 2005-01-18, http://yaml.org/type/timestamp.html

84 1.7 Non-Normative References

[Apache]	Apache Server, https://httpd.apache.org/
[Chef]	Chef, https://wiki.opscode.com/display/chef/Home
[NodeJS]	Node.js, https://nodejs.org/
[Puppet]	Puppet, http://puppetlabs.com/
[WordPress]	WordPress, https://wordpress.org/
[Maven-Version]	Apache Maven version policy draft: https://cwiki.apache.org/confluence/display/MAVEN/Version+number+policy

85 1.8 Glossary

86 The following terms are used throughout this specification and have the following definitions when used in
87 context of this document.

Term	Definition
Instance Model	A deployed service is a running instance of a Service Template. More precisely, the instance is derived by instantiating the Topology Template of its Service Template, most often by running a special plan defined for the Service Template, often referred to as build plan.
Node Template	A <i>Node Template</i> specifies the occurrence of a software component node as part of a Topology Template. Each Node Template refers to a Node Type that defines the semantics of the node (e.g., properties, attributes, requirements, capabilities, interfaces). Node Types are defined separately for reuse purposes.
Relationship Template	A <i>Relationship Template</i> specifies the occurrence of a relationship between nodes in a Topology Template. Each Relationship Template refers to a Relationship Type that defines the semantics relationship (e.g., properties, attributes, interfaces, etc.). Relationship Types are defined separately for reuse purposes.
Service Template	A <i>Service Template</i> is typically used to specify the “topology” (or structure) and “orchestration” (or invocation of management behavior) of IT services so that they can be provisioned and managed in accordance with constraints and policies. Specifically, TOSCA Service Templates optionally allow definitions of a TOSCA Topology Template , TOSCA types (e.g., Node, Relationship, Capability, Artifact, etc.), groupings, policies and constraints along with any input or output declarations.
Topology Model	The term Topology Model is often used synonymously with the term Topology Template with the use of “model” being prevalent when considering a Service Template’s topology definition as an abstract representation of an application or service to facilitate understanding of its functional components and by eliminating unnecessary details.
Topology Template	A Topology Template defines the structure of a service in the context of a Service

Template. A Topology Template consists of a set of Node Template and Relationship Template definitions that together define the topology model of a service as a (not necessarily connected) directed graph.

The term Topology Template is often used synonymously with the term [Topology Model](#). The distinction is that a topology template can be used to instantiate and orchestrate the model as a **reusable pattern** and includes all details necessary to accomplish it.

Abstract Node Template	An abstract node template is a node that doesn't define an implementation artifact for the create operation of the TOSCA lifecycle. The create operation can be delegated to the TOSCA Orchestrator. Being delegated an abstract node may not be able to execute user provided implementation artifacts for operations post create (for example configure, start etc.).
No-Op Node Template	A No-Op node template is a specific abstract node template that does not specify any implementation for any operation.

88 2 TOSCA by example

89 This **non-normative** section contains several sections that show how to model applications with TOSCA
90 Simple Profile using YAML by example starting with a “Hello World” template up through examples that
91 show complex composition modeling.

92 2.1 A “hello world” template for TOSCA Simple Profile in YAML

93 As mentioned before, the TOSCA simple profile assumes the existence of a small set of pre-defined,
94 normative set of node types (e.g., a ‘Compute’ node) along with other types, which will be introduced
95 through the course of this document, for creating TOSCA Service Templates. It is envisioned that many
96 additional node types for building service templates will be created by communities some may be
97 published as profiles that build upon the TOSCA Simple Profile specification. Using the normative TOSCA
98 Compute node type, a very basic “Hello World” TOSCA template for deploying just a single server would
99 look as follows:

100 Example 1 - TOSCA Simple "Hello World"

```
tosca_definitions_version: tosca_simple_yaml_1_0

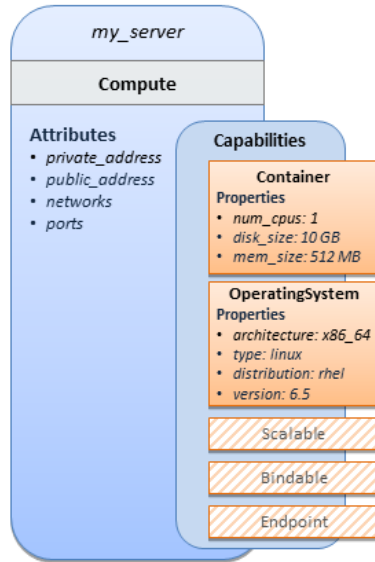
description: Template for deploying a single server with predefined properties.

topology_template:
  node_templates:
    my_server:
      type: tosca.nodes.Compute
      capabilities:
        # Host container properties
        host:
          properties:
            num_cpus: 1
            disk_size: 10 GB
            mem_size: 4096 MB
        # Guest Operating System properties
        os:
          properties:
            # host Operating System image properties
            architecture: x86_64
            type: linux
            distribution: rhel
            version: 6.5
```

101 The template above contains a very simple topology template with only a single ‘Compute’ node template
102 that declares some basic values for properties within two of the several capabilities that are built into the
103 Compute node type definition. All TOSCA Orchestrators are expected to know how to instantiate a
104 Compute node since it is normative and expected to represent a well-known function that is portable
105 across TOSCA implementations. This expectation is true for all normative TOSCA Node and
106 Relationship types that are defined in the Simple Profile specification. This means, with TOSCA’s
107 approach, that the application developer does not need to provide any deployment or implementation
108 artifacts that contain code or logic to orchestrate these common software components. TOSCA
109 orchestrators simply select or allocate the correct node (resource) type that fulfills the application
110 topologies requirements using the properties declared in the node and its capabilities.

111 In the above example, the “**host**” capability contains properties that allow application developers to
112 optionally supply the number of CPUs, memory size and disk size they believe they need when the

113 Compute node is instantiated in order to run their applications. Similarly, the “os” capability is used to
 114 provide values to indicate what host operating system the Compute node should have when it is
 115 instantiated.
 116
 117 The logical diagram of the “hello world” Compute node would look as follows:



118
 119
 120 As you can see, the **Compute** node also has attributes and other built-in capabilities, such as **Bindable**
 121 and **Endpoint**, each with additional properties that will be discussed in other examples later in this
 122 document. Although the Compute node has no direct properties apart from those in its capabilities, other
 123 TOSCA node type definitions may have properties that are part of the node type itself in addition to
 124 having Capabilities. TOSCA orchestration engines are expected to validate all property values provided
 125 in a node template against the property definitions in their respective node type definitions referenced in
 126 the service template. The **tosca_definitions_version** keyname in the TOSCA service template
 127 identifies the versioned set of normative TOSCA type definitions to use for validating those types defined
 128 in the TOSCA Simple Profile including the Compute node type. Specifically, the value
 129 **tosca_simple_yaml_1_0** indicates Simple Profile v1.0.0 definitions would be used for validation. Other
 130 type definitions may be imported from other service templates using the **import** keyword discussed later.

131 **2.1.1 Requesting input parameters and providing output**

132 Typically, one would want to allow users to customize deployments by providing input parameters instead
 133 of using hardcoded values inside a template. In addition, output values are provided to pass information
 134 that perhaps describes the state of the deployed template to the user who deployed it (such as the private
 135 IP address of the deployed server). A refined service template with corresponding **inputs** and **outputs**
 136 sections is shown below.

137 **Example 2 - Template with input and output parameter sections**

```
tosca_definitions_version: toscasimple_yaml_1_0

description: Template for deploying a single server with predefined properties.

topology_template:
  inputs:
    cpus:
      type: integer
```

```

description: Number of CPUs for the server.
constraints:
  - valid_values: [ 1, 2, 4, 8 ]

node_templates:
  my_server:
    type: tosca.nodes.Compute
    capabilities:
      # Host container properties
      host:
        properties:
          # Compute properties
          num_cpus: { get_input: cpus }
          mem_size: 2048 MB
          disk_size: 10 GB

    outputs:
      server_ip:
        description: The private IP address of the provisioned server.
        value: { get_attribute: [ my_server, private_address ] }

```

138 The **inputs** and **outputs** sections are contained in the **topology_template** element of the TOSCA
 139 template, meaning that they are scoped to node templates within the topology template. Input parameters
 140 defined in the inputs section can be assigned to properties of node template within the containing
 141 topology template; output parameters can be obtained from attributes of node templates within the
 142 containing topology template.

143 Note that the **inputs** section of a TOSCA template allows for defining optional constraints on each input
 144 parameter to restrict possible user input. Further note that TOSCA provides for a set of intrinsic functions
 145 like **get_input**, **get_property** or **get_attribute** to reference elements within the template or to
 146 retrieve runtime values.

147 2.2 TOSCA template for a simple software installation

148 Software installations can be modeled in TOSCA as node templates that get related to the node template
 149 for a server on which the software would be installed. With a number of existing software node types (e.g.
 150 either created by the TOSCA work group or a community) template authors can just use those node types
 151 for writing service templates as shown below.

152 Example 3 - Simple (MySQL) software installation on a TOSCA Compute node

```

tosca_definitions_version: tosca_simple_yaml_1_0
description: Template for deploying a single server with MySQL software on top.

topology_template:
  inputs:
    # omitted here for brevity

  node_templates:
    mysql:
      type: tosca.nodes.DBMS.MySQL
      properties:
        root_password: { get_input: my_mysql_rootpw }
        port: { get_input: my_mysql_port }
      requirements:
        - host: db_server

```

```

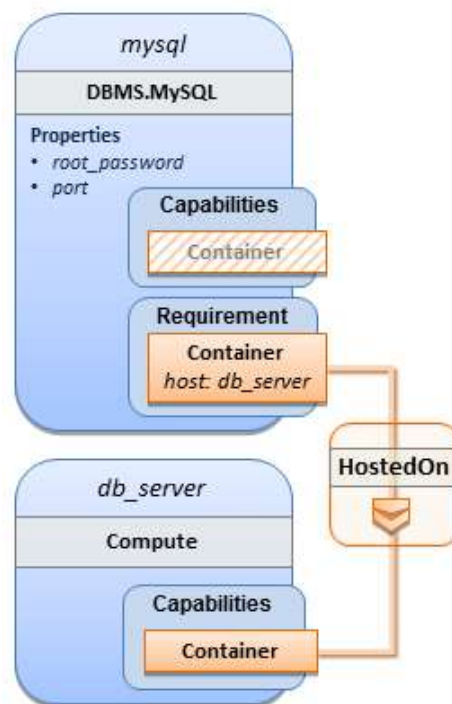
db_server:
  type: tosca.nodes.Compute
  capabilities:
    # omitted here for brevity

```

153 The example above makes use of a node type `tosca.nodes.DBMS.MySQL` for the `mysql` node template to
 154 install MySQL on a server. This node type allows for setting a property `root_password` to adapt the
 155 password of the MySQL root user at deployment. The set of properties and their schema has been
 156 defined in the node type definition. By means of the `get_input` function, a value provided by the user at
 157 deployment time is used as value for the `root_password` property. The same is true for the `port`
 158 property.

159 The `mysql` node template is related to the `db_server` node template (of type `tosca.nodes.Compute`) via
 160 the `requirements` section to indicate where MySQL is to be installed. In the TOSCA metamodel, nodes
 161 get related to each other when one node has a requirement against some feature provided by another
 162 node. What kinds of requirements exist is defined by the respective node type. In case of MySQL, which
 163 is software that needs to be installed or hosted on a compute resource, the underlying node type named
 164 `DBMS` has a predefined requirement called `host`, which needs to be fulfilled by pointing to a node template
 165 of type `tosca.nodes.Compute`.

166 The logical relationship between the `mysql` node and its host `db_server` node would appear as follows:



167

168 Within the `requirements` section, all entries simple entries are a map which contains the symbolic name
 169 of a requirement definition as the `key` and the identifier of the fulfilling node as the `value`. The value is
 170 essentially the symbolic name of the other node template; specifically, or the example above, the `host`
 171 requirement is fulfilled by referencing the `db_server` node template. The underlying TOSCA `DBMS` node
 172 type already defines a complete requirement definition for the `host` requirement of type `Container` and
 173 assures that a `HostedOn` TOSCA relationship will automatically be created and will only allow a valid
 174 target host node is of type `Compute`. This approach allows the template author to simply provide the
 175 name of a valid `Compute` node (i.e., `db_server`) as the value for the `mysql` node's `host` requirement and
 176 not worry about defining anything more complex if they do not want to.

177 2.3 Overriding behavior of predefined node types

178 Node types in TOSCA have associated implementations that provide the automation (e.g. in the form of
179 scripts such as Bash, Chef or Python) for the normative lifecycle operations of a node. For example, the
180 node type implementation for a MySQL database would associate scripts to TOSCA node operations like
181 **configure**, **start**, or **stop** to manage the state of MySQL at runtime.

182 Many node types may already come with a set of operational scripts that contain basic commands that
183 can manage the state of that specific node. If it is desired, template authors can provide a custom script
184 for one or more of the operation defined by a node type in their node template which will override the
185 default implementation in the type. The following example shows a **mysql** node template where the
186 template author provides their own configure script:

187 Example 4 - Node Template overriding its Node Type's "configure" interface

```
tosca_definitions_version: tosca_simple_yaml_1_0

description: Template for deploying a single server with MySQL software on top.

topology_template:
  inputs:
    # omitted here for brevity

  node_templates:
    mysql:
      type: tosca.nodes.DBMS.MySQL
      properties:
        root_password: { get_input: my_mysql_rootpw }
        port: { get_input: my_mysql_port }
      requirements:
        - host: db_server
      interfaces:
        Standard:
          configure: scripts/my_own_configure.sh

    db_server:
      type: tosca.nodes.Compute
      capabilities:
        # omitted here for brevity
```

188 In the example above, the **my_own_configure.sh** script is provided for the **configure** operation of the
189 MySQL node type's **Standard** lifecycle interface. The path given in the example above (i.e., 'scripts/') is
190 interpreted relative to the template file, but it would also be possible to provide an absolute URI to the
191 location of the script.

192 In other words, operations defined by node types can be thought of as "hooks" into which automation can
193 be injected. Typically, node type implementations provide the automation for those "hooks". However,
194 within a template, custom automation can be injected to run in a hook in the context of the one, specific
195 node template (i.e. without changing the node type).

196 2.4 TOSCA template for database content deployment

197 In the Example 4, shown above, the deployment of the MySQL middleware only, i.e. without actual
198 database content was shown. The following example shows how such a template can be extended to
199 also contain the definition of custom database content on-top of the MySQL DBMS software.


```

tosca_definitions_version: tosca_simple_yaml_1_0

description: Template for deploying MySQL and database content.

topology_template:
  inputs:
    # omitted here for brevity

  node_templates:
    my_db:
      type: tosca.nodes.Database.MySQL
      properties:
        name: { get_input: database_name }
        user: { get_input: database_user }
        password: { get_input: database_password }
        port: { get_input: database_port }
      artifacts:
        db_content:
          file: files/my_db_content.txt
          type: tosca.artifacts.File
      requirements:
        - host: mysql
      interfaces:
        Standard:
          create:
            implementation: db_create.sh
            inputs:
              # Copy DB file artifact to server's staging area
              db_data: { get_artifact: [ SELF, db_content ] }

    mysql:
      type: tosca.nodes.DBMS.MySQL
      properties:
        root_password: { get_input: mysql_rootpw }
        port: { get_input: mysql_port }
      requirements:
        - host: db_server

    db_server:
      type: tosca.nodes.Compute
      capabilities:
        # omitted here for brevity

```

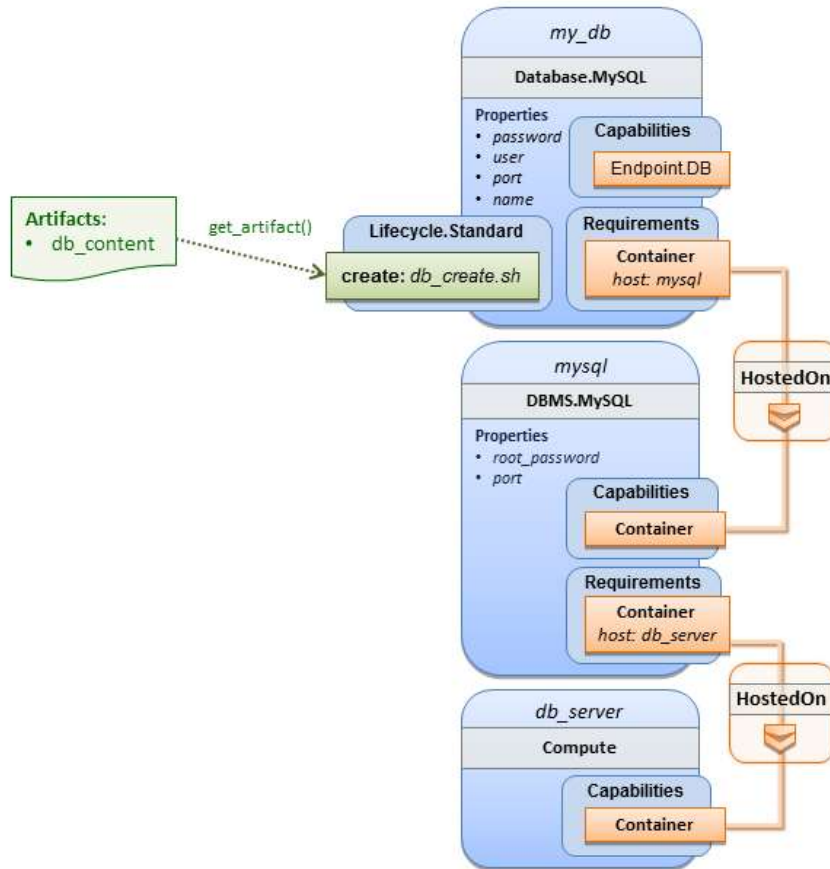
201 In the example above, the **my_db** node template or type **tosca.nodes.Database.MySQL** represents an
 202 actual MySQL database instance managed by a MySQL DBMS installation. The **requirements** section of
 203 the **my_db** node template expresses that the database it represents is to be hosted on a MySQL DBMS
 204 node template named **mysql** which is also declared in this template.

205 In its **artifacts** section of the **my_db** the node template, there is an artifact definition named **db_content**
 206 which represents a text file **my_db_content.txt** which in turn will be used to add content to the SQL
 207 database as part of the **create** operation. The **requirements** section of the **my_db** node template
 208 expresses that the database is hosted on a MySQL DBMS represented by the **mysql** node.

209 As you can see above, a script is associated with the create operation with the name **db_create.sh**.
 210 The TOSCA Orchestrator sees that this is not a named artifact declared in the node's artifact section, but

211 instead a filename for a normative TOSCA implementation artifact script type (i.e.,
 212 **tosca.artifacts.Implementation.Bash**). Since this is an implementation type for TOSCA, the
 213 orchestrator will execute the script automatically to create the node on **db_server**, but first it will prepare
 214 the local environment with the declared inputs for the operation. In this case, the orchestrator would see
 215 that the **db_data** input is using the **get_artifact** function to retrieve the file (**my_db_content.txt**)
 216 which is associated with the **db_content** artifact name prior to executing the **db_create.sh** script.

217 The logical diagram for this example would appear as follows:



218
 219 Note that while it would be possible to define one node type and corresponding node templates that
 220 represent both the DBMS middleware and actual database content as one entity, TOSCA normative node
 221 types distinguish between middleware (container) and application (containeer) node types. This allows on
 222 one hand to have better re-use of generic middleware node types without binding them to content running
 223 on top of them, and on the other hand this allows for better substitutability of, for example, middleware
 224 components like a DBMS during the deployment of TOSCA models.

225 2.5 TOSCA template for a two-tier application

226 The definition of multi-tier applications in TOSCA is quite similar to the example shown in section 2.2, with
 227 the only difference that multiple software node stacks (i.e., node templates for middleware and application
 228 layer components), typically hosted on different servers, are defined and related to each other. The
 229 example below defines a web application stack hosted on the **web_server** “compute” resource, and a
 230 database software stack similar to the one shown earlier in section 6 hosted on the **db_server** compute
 231 resource.

232 Example 6 - Basic two-tier application (web application and database server tiers)

```
tosca_definitions_version: tosca_simple_yaml_1_0
```

description: Template for deploying a two-tier application servers on two

topology_template:

inputs:

Admin user name and password to use with the WordPress application

wp_admin_username:

type: string

wp_admin_password:

type: string

wp_db_name:

type: string

wp_db_user:

type: string

wp_db_password:

type: string

wp_db_port:

type: integer

mysql_root_password:

type: string

mysql_port:

type: integer

context_root:

type: string

node_templates:

wordpress:

type: toska.nodes.WebApplication.WordPress

properties:

context_root: { get_input: context_root }

admin_user: { get_input: wp_admin_username }

admin_password: { get_input: wp_admin_password }

db_host: { get_attribute: [db_server, private_address] }

requirements:

- host: apache

- database_endpoint: wordpress_db

interfaces:

Standard:

inputs:

db_host: { get_attribute: [db_server, private_address] }

db_port: { get_property: [wordpress_db, port] }

db_name: { get_property: [wordpress_db, name] }

db_user: { get_property: [wordpress_db, user] }

db_password: { get_property: [wordpress_db, password] }

apache:

type: toska.nodes.WebServer.Apache

properties:

omitted here for brevity

requirements:

- host: web_server

web_server:

type: toska.nodes.Compute

capabilities:

```

# omitted here for brevity

wordpress_db:
  type: tosca.nodes.Database.MySQL
  properties:
    name: { get_input: wp_db_name }
    user: { get_input: wp_db_user }
    password: { get_input: wp_db_password }
    port: { get_input: wp_db_port }
  requirements:
    - host: mysql

mysql:
  type: tosca.nodes.DBMS.MySQL
  properties:
    root_password: { get_input: mysql_root_password }
    port: { get_input: mysql_port }
  requirements:
    - host: db_server

db_server:
  type: tosca.nodes.Compute
  capabilities:
    # omitted here for brevity

```

233 The web application stack consists of the **wordpress** [WordPress], the **apache** [Apache] and the
 234 **web_server** node templates. The **wordpress** node template represents a custom web application of type
 235 **tosca.nodes.WebApplication.WordPress** which is hosted on an Apache web server represented by the
 236 **apache** node template. This hosting relationship is expressed via the **host** entry in the **requirements**
 237 section of the **wordpress** node template. The **apache** node template, finally, is hosted on the
 238 **web_server** compute node.

239 The database stack consists of the **wordpress_db**, the **mysql** and the **db_server** node templates. The
 240 **wordpress_db** node represents a custom database of type **tosca.nodes.Database.MySQL** which is
 241 hosted on a MySQL DBMS represented by the **mysql** node template. This node, in turn, is hosted on the
 242 **db_server** compute node.

243 The **wordpress** node requires a connection to the **wordpress_db** node, since the WordPress application
 244 needs a database to store its data in. This relationship is established through the **database_endpoint**
 245 entry in the **requirements** section of the **wordpress** node template's declared node type. For configuring
 246 the WordPress web application, information about the database to connect to is required as input to the
 247 **configure** operation. Therefore, the input parameters are defined and values for them are retrieved from
 248 the properties and attributes of the **wordpress_db** node via the **get_property** and **get_attribute**
 249 functions. In the above example, these inputs are defined at the interface-level and would be available to
 250 all operations of the **Standard** interface (i.e., the **tosca.interfaces.node.lifecycle.Standard**
 251 interface) within the **wordpress** node template and not just the **configure** operation.

252 2.6 Using a custom script to establish a relationship in a template

253 In previous examples, the template author did not have to think about explicit relationship types to be
 254 used to link a requirement of a node to another node of a model, nor did the template author have to think
 255 about special logic to establish those links. For example, the **host** requirement in previous examples just
 256 pointed to another node template and based on metadata in the corresponding node type definition the
 257 relationship type to be established is implicitly given.

258 In some cases, it might be necessary to provide special processing logic to be executed when
 259 establishing relationships between nodes at runtime. For example, when connecting the WordPress

260 application from previous examples to the MySQL database, it might be desired to apply custom
261 configuration logic in addition to that already implemented in the application node type. In such a case, it
262 is possible for the template author to provide a custom script as implementation for an operation to be
263 executed at runtime as shown in the following example.

264 **Example 7 - Providing a custom relationship script to establish a connection**

```
tosca_definitions_version: tosca_simple_yaml1_1_0

description: Template for deploying a two-tier application on two servers.

topology_template:
  inputs:
    # omitted here for brevity

  node_templates:
    wordpress:
      type: tosca.nodes.WebApplication.WordPress
      properties:
        # omitted here for brevity
      requirements:
        - host: apache
        - database_endpoint:
            node: wordpress_db
            relationship: my_custom_database_connection

    wordpress_db:
      type: tosca.nodes.Database.MySQL
      properties:
        # omitted here for the brevity
      requirements:
        - host: mysql

  relationship_templates:
    my_custom_database_connection:
      type: ConnectsTo
      interfaces:
        Configure:
          pre_configure_source: scripts/wp_db_configure.sh

# other resources not shown for this example ...
```

265 The node type definition for the **wordpress** node template is **WordPress** which declares the complete
266 **database_endpoint** requirement definition. This **database_endpoint** declaration indicates it must be
267 fulfilled by any node template that provides an **Endpoint.Database** Capability Type using a **ConnectsTo**
268 relationship. The **wordpress_db** node template's underlying **MySQL** type definition indeed provides the
269 **Endpoint.Database** Capability type. In this example however, no explicit relationship template is
270 declared; therefore, TOSCA orchestrators would automatically create a **ConnectsTo** relationship to
271 establish the link between the **wordpress** node and the **wordpress_db** node at runtime.

272 The **ConnectsTo** relationship (see 5.7.4) also provides a default **Configure** interface with operations that
273 optionally get executed when the orchestrator establishes the relationship. In the above example, the
274 author has provided the custom script **wp_db_configure.sh** to be executed for the operation called
275 **pre_configure_source**. The script file is assumed to be located relative to the referencing service
276 template such as a relative directory within the TOSCA Cloud Service Archive (CSAR) packaging format.
277 This approach allows for conveniently hooking in custom behavior without having to define a completely
278 new derived relationship type.

279 2.7 Using custom relationship types in a TOSCA template

280 In the previous section it was shown how custom behavior can be injected by specifying scripts inline in
281 the requirements section of node templates. When the same custom behavior is required in many
282 templates, it does make sense to define a new relationship type that encapsulates the custom behavior in
283 a re-usable way instead of repeating the same reference to a script (or even references to multiple
284 scripts) in many places.

285 Such a custom relationship type can then be used in templates as shown in the following example.

286 Example 8 - A web application Node Template requiring a custom database connection type

```
tosca_definitions_version: tosca_simple_yaml_1_0

description: Template for deploying a two-tier application on two servers.

topology_template:
  inputs:
    # omitted here for brevity

  node_templates:
    wordpress:
      type: tosca.nodes.WebApplication.WordPress
      properties:
        # omitted here for brevity
      requirements:
        - host: apache
        - database_endpoint:
            node: wordpress_db
            relationship: my.types.WordpressDbConnection

    wordpress_db:
      type: tosca.nodes.Database.MySQL
      properties:
        # omitted here for the brevity
      requirements:
        - host: mysql

# other resources not shown here ...
```

287 In the example above, a special relationship type **my.types.WordpressDbConnection** is specified for
288 establishing the link between the **wordpress** node and the **wordpress_db** node through the use of the
289 **relationship** (keyword) attribute in the **database** reference. It is assumed, that this special relationship
290 type provides some extra behavior (e.g., an operation with a script) in addition to what a generic
291 “connects to” relationship would provide. The definition of this custom relationship type is shown in the
292 following section.

293 2.7.1 Definition of a custom relationship type

294 The following YAML snippet shows the definition of the custom relationship type used in the previous
295 section. This type derives from the base “ConnectsTo” and overrides one operation defined by that base
296 relationship type. For the **pre_configure_source** operation defined in the **Configure** interface of the
297 ConnectsTo relationship type, a script implementation is provided. It is again assumed that the custom
298 configure script is located at a location relative to the referencing service template, perhaps provided in
299 some application packaging format (e.g., the TOSCA Cloud Service Archive (CSAR) format).

300 Example 9 - Defining a custom relationship type

```
tosca_definitions_version: tosca_simple_yaml_1_0

description: Definition of custom WordpressDbConnection relationship type

relationship_types:
  my.types.WordpressDbConnection:
    derived_from: tosca.relationships.ConnectsTo
    interfaces:
      Configure:
        pre_configure_source: scripts/wp_db_configure.sh
```

301 In the above example, the **Configure** interface is the specified alias or shorthand name for the TOSCA
302 interface type with the full name of **tosca.interfaces.relationship.Configure** which is defined in
303 the appendix.

304 2.8 Defining generic dependencies between nodes in a template

305 In some cases, it can be necessary to define a generic dependency between two nodes in a template to
306 influence orchestration behavior, i.e. to first have one node processed before another dependent node
307 gets processed. This can be done by using the generic **dependency** requirement which is defined by the
308 [TOSCA Root Node Type](#) and thus gets inherited by all other node types in TOSCA (see section 5.9.1).

309 Example 10 - Simple dependency relationship between two nodes

```
tosca_definitions_version: tosca_simple_yaml_1_0

description: Template with a generic dependency between two nodes.

topology_template:
  inputs:
    # omitted here for brevity

  node_templates:
    my_app:
      type: my.types.MyApplication
      properties:
        # omitted here for brevity
      requirements:
        - dependency: some_service

    some_service:
      type: some.nodetype.SomeService
      properties:
        # omitted here for brevity
```

310 As in previous examples, the relation that one node depends on another node is expressed in the
311 **requirements** section using the built-in requirement named **dependency** that exists for all node types in
312 TOSCA. Even if the creator of the **MyApplication** node type did not define a specific requirement for
313 **SomeService** (similar to the **database** requirement in the example in section 2.6), the template author
314 who knows that there is a timing dependency and can use the generic **dependency** requirement to
315 express that constraint using the very same syntax as used for all other references.

316 2.9 Describing abstract requirements for nodes and capabilities in a 317 TOSCA template

318 In TOSCA templates, nodes are either:

- 319 • **Concrete:** meaning that they have a deployment and/or one or more implementation artifacts that
320 are declared on the “create” operation of the node’s Standard lifecycle interface, or they are
 - 321 • **Abstract:** where the template describes the node type along with its required capabilities and
322 properties that must be satisfied.
- 323

324 TOSCA Orchestrators, by default, when finding an abstract node in TOSCA Service Template during
325 deployment will attempt to “select” a concrete implementation for the abstract node type that best
326 matches and fulfills the requirements and property constraints the template author provided for that
327 abstract node. The concrete implementation of the node could be provided by another TOSCA Service
328 Template (perhaps located in a catalog or repository known to the TOSCA Orchestrator) or by an existing
329 resource or service available within the target Cloud Provider’s platform that the TOSCA Orchestrator
330 already has knowledge of.

331
332 TOSCA supports two methods for template authors to express requirements for an abstract node within a
333 TOSCA service template.

- 334
335 1. **Using a target node filter:** where a node template can describe a requirement (relationship) for
336 another node without including it in the topology. Instead, the node provides a `node_filter` to
337 describe the target node type along with its capabilities and property constraints
- 338
339 2. **Using an abstract node template:** that describes the abstract node’s type along with its property
340 constraints and any requirements and capabilities it also exports. This first method you have
341 already seen in examples from previous chapters where the Compute node is abstract and
342 selectable by the TOSCA Orchestrator using the supplied Container and [OperatingSystem](#)
343 capabilities property constraints.

344
345 These approaches allow architects and developers to create TOSCA service templates that are
346 composable and can be reused by allowing flexible matching of one template’s requirements to another’s
347 capabilities. Examples of both these approaches are shown below.

348 2.9.1 Using a node_filter to define hosting infrastructure requirements for a 349 software

350 Using TOSCA, it is possible to define only the software components of an application in a template and
351 just express constrained requirements against the hosting infrastructure. At deployment time, the provider
352 can then do a late binding and dynamically allocate or assign the required hosting infrastructure and
353 place software components on top.

354 This example shows how a single software component (i.e., the mysql node template) can define its **host**
355 requirements that the TOSCA Orchestrator and provider will use to select or allocate an appropriate host
356 **Compute** node by using matching criteria provided on a **node_filter**.

357 Example 11 - An abstract "host" requirement using a node filter

```
tosca_definitions_version: tosca_simple_yaml_1_0  
  
description: Template with requirements against hosting infrastructure.
```



```

topology_template:
  inputs:
    # omitted here for brevity

  node_templates:
    mysql:
      type: tosca.nodes.DBMS.MySQL
      properties:
        # omitted here for brevity
      requirements:
        - host:
            node_filter:
              capabilities:
                # Constraints for selecting "host" (Container Capability)
                - host:
                    properties:
                      - num_cpus: { in_range: [ 1, 4 ] }
                      - mem_size: { greater_or_equal: 2 GB }
                # Constraints for selecting "os" (OperatingSystem Capability)
                - os:
                    properties:
                      - architecture: { equal: x86_64 }
                      - type: linux
                      - distribution: ubuntu

```

358 In the example above, the `mysql` component contains a `host` requirement for a node of type `Compute`
 359 which it inherits from its parent DBMS node type definition; however, there is no declaration or reference
 360 to any node template of type `Compute`. Instead, the `mysql` node template augments the abstract “`host`”
 361 requirement with a `node_filter` which contains additional selection criteria (in the form of property
 362 constraints that the provider must use when selecting or allocating a host `Compute` node.

363 Some of the constraints shown above narrow down the boundaries of allowed values for certain
 364 properties such as `mem_size` or `num_cpus` for the “`host`” capability by means of qualifier functions such
 365 as `greater_or_equal`. Other constraints, express specific values such as for the `architecture` or
 366 `distribution` properties of the “`os`” capability which will require the provider to find a precise match.

367 Note that when no qualifier function is provided for a property (filter), such as for the `distribution`
 368 property, it is interpreted to mean the `equal` operator as shown on the `architecture` property.

369 2.9.2 Using an abstract node template to define infrastructure requirements 370 for software

371 This previous approach works well if no other component (i.e., another node template) other than `mysql`
 372 node template wants to reference the same `Compute` node the orchestrator would instantiate. However,
 373 perhaps another component wants to also be deployed on the same host, yet still allow the flexible
 374 matching achieved using a node-filter. The alternative to the above approach is to create an abstract
 375 node template that represents the `Compute` node in the topology as follows:

376 Example 12 - An abstract Compute node template with a node filter

```

tosca_definitions_version: tosca_simple_yaml_1_0

description: Template with requirements against hosting infrastructure.

topology_template:

```

```

inputs:
  # omitted here for brevity

node_templates:
  mysql:
    type: toska.nodes.DBMS.MySQL
    properties:
      # omitted here for brevity
    requirements:
      - host: mysql_compute

# Abstract node template (placeholder) to be selected by provider
mysql_compute:
  type: Compute
  node_filter:
    capabilities:
      - host:
          properties:
            num_cpus: { equal: 2 }
            mem_size: { greater_or_equal: 2 GB }
      - os:
          properties:
            architecture: { equal: x86_64 }
            type: linux
            distribution: ubuntu

```

377 As you can see the resulting `mysql_compute` node template looks very much like the “hello world”
 378 template as shown in [Chapter 2.1](#) (where the `Compute` node template was abstract), but this one also
 379 allows the TOSCA orchestrator more flexibility when “selecting” a host `Compute` node by providing flexible
 380 constraints for properties like `mem_size`.

381 As we proceed, you will see that TOSCA provides many normative node types like `Compute` for
 382 commonly found services (e.g., `BlockStorage`, `WebServer`, `Network`, etc.). When these TOSCA
 383 normative node types are used in your application’s topology they are always assumed to be “selectable”
 384 by TOSCA Orchestrators which work with target infrastructure providers to find or allocate the best match
 385 for them based upon your application’s requirements and constraints.

386 2.9.3 Using a `node_filter` to define requirements on a database for an 387 application

388 In the same way requirements can be defined on the hosting infrastructure (as shown above) for an
 389 application, it is possible to express requirements against application or middleware components such as
 390 a database that is not defined in the same template. The provider may then allocate a database by any
 391 means, (e.g. using a database-as-a-service solution).

392 **Example 13 - An abstract database requirement using a node filter**

```

tosca_definitions_version: toska_simple_yaml_1_0

description: Template with a TOSCA Orchestrator selectable database requirement
using a node_filter.

topology_template:
  inputs:

```

```

# omitted here for brevity

node_templates:
  my_app:
    type: my.types.MyApplication
    properties:
      admin_user: { get_input: admin_username }
      admin_password: { get_input: admin_password }
      db_endpoint_url: { get_property: [SELF, database_endpoint, url_path ] }
    requirements:
      - database_endpoint:
          node: my.types.nodes.MyDatabase
          node_filter:
            properties:
              - db_version: { greater_or_equal: 5.5 }

```

393 In the example above, the application **my_app** requires a database node of type **MyDatabase** which has a
 394 **db_version** property value of **greater_or_equal** to the value 5.5.

395 This example also shows how the **get_property** intrinsic function can be used to retrieve the **url_path**
 396 property from the database node that will be selected by the provider and connected to **my_app** at runtime
 397 due to fulfillment of the **database_endpoint** requirement. To locate the property, the **get_property**'s first
 398 argument is set to the keyword **SELF** which indicates the property is being referenced from something in
 399 the node itself. The second parameter is the name of the requirement named **database_endpoint** which
 400 contains the property we are looking for. The last argument is the name of the property itself (i.e.,
 401 **url_path**) which contains the value we want to retrieve and assign to **db_endpoint_url**.

402 The alternative representation, which includes a node template in the topology for database that is still
 403 selectable by the TOSCA orchestrator for the above example, is as follows:

404 **Example 14 - An abstract database node template**

```

tosca_definitions_version: tosca_simple_yaml_1_0

description: Template with a TOSCA Orchestrator selectable database using node
template.

topology_template:
  inputs:
    # omitted here for brevity

  node_templates:
    my_app:
      type: my.types.MyApplication
      properties:
        admin_user: { get_input: admin_username }
        admin_password: { get_input: admin_password }
        db_endpoint_url: { get_property: [SELF, database_endpoint, url_path ] }
      requirements:
        - database_endpoint: my_abstract_database

    my_abstract_database:
      type: my.types.nodes.MyDatabase
      properties:
        - db_version: { greater_or_equal: 5.5 }

```

405 **2.10 Using node template substitution for model composition**

406 From an application perspective, it is often not necessary or desired to dive into platform details, but the
407 platform/runtime for an application is abstracted. In such cases, the template for an application can use
408 generic representations of platform components. The details for such platform components, such as the
409 underlying hosting infrastructure at its configuration, can then be defined in separate template files that
410 can be used for substituting the more abstract representations in the application level template file.

411 **2.10.1 Understanding node template instantiation through a TOSCA** 412 **Orchestrator**

413 When a topology template is instantiated by a TOSCA Orchestrator, the orchestrator has to look for
414 realizations of the single node templates according to the node types specified for each node template.
415 Such realizations can either be node types that include the appropriate implementation artifacts and
416 deployment artifacts that can be used by the orchestrator to bring to life the real-world resource modeled
417 by a node template. Alternatively, separate topology templates may be annotated as being suitable for
418 realizing a node template in the top-level topology template.

419

420 In the latter case, a TOSCA Orchestrator will use additional substitution mapping information provided as
421 part of the substituting topology templates to derive how the substituted part gets “wired” into the overall
422 deployment, for example, how capabilities of a node template in the top-level topology template get
423 bound to capabilities of node templates in the substituting topology template.

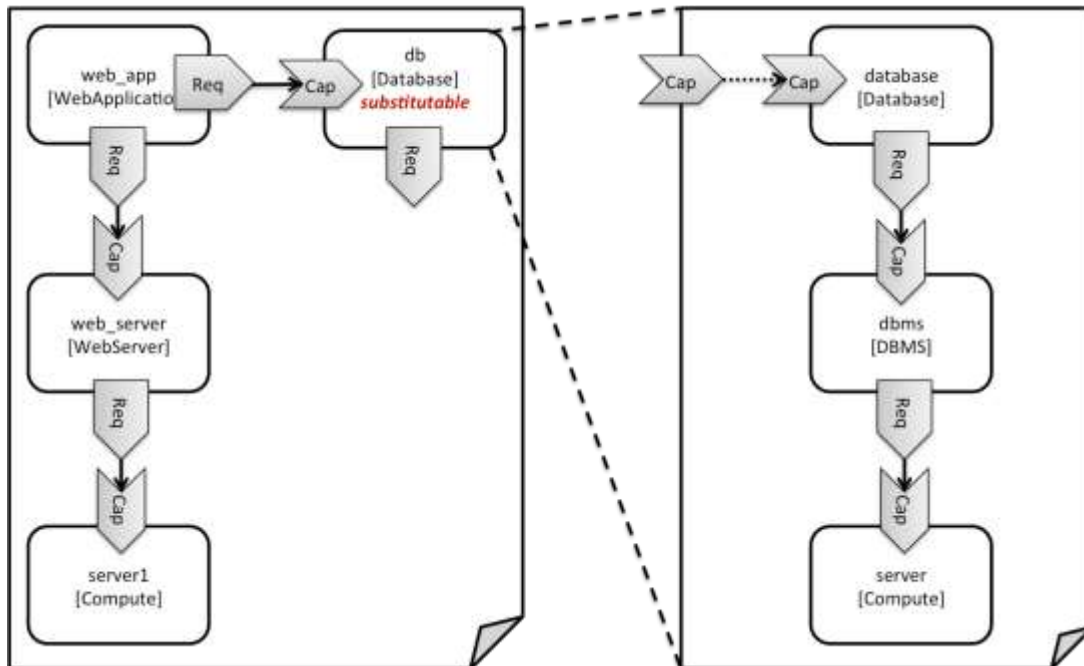
424

425 Thus, in cases where no “normal” node type implementation is available, or the node type corresponds to
426 a whole subsystem that cannot be implemented as a single node, additional topology templates can be
427 used for filling in more abstract placeholders in top level application templates.

428 **2.10.2 Definition of the top-level service template**

429 The following sample defines a web application **web_app** connected to a database **db**. In this example, the
430 complete hosting stack for the application is defined within the same topology template: the web
431 application is hosted on a web server **web_server**, which in turn is installed (hosted) on a compute node
432 **server**.

433 The hosting stack for the database **db**, in contrast, is not defined within the same file but only the
434 database is represented as a node template of type **tosca.nodes.Database**. The underlying hosting
435 stack for the database is defined in a separate template file, which is shown later in this section. Within
436 the current template, only a number of properties (**user**, **password**, **name**) are assigned to the database
437 using hardcoded values in this simple example.



438

439

Figure 1: Using template substitution to implement a database tier

440 When a node template is to be substituted by another service template, this has to be indicated to an
 441 orchestrator by means of a special “*substitutable*” directive. This directive causes, for example, special
 442 processing behavior when validating the left-hand service template in Figure 1. The hosting requirement
 443 of the **db** node template is not bound to any capability defined within the service template, which would
 444 normally cause a validation error. When the “*substitutable*” directive is present, the orchestrator will
 445 however first try to perform substitution of the respective node template and after that validate if all
 446 mandatory requirements of all nodes in the resulting graph are fulfilled.

447

448 Note that in contrast to the use case described in section 2.9.2 (where a database was abstractly referred
 449 to in the **requirements** section of a node and the database itself was not represented as a node
 450 template), the approach shown here allows for some additional modeling capabilities in cases where this
 451 is required.

452

453 For example, if multiple components need to use the same database (or any other sub-system of the
 454 overall service), this can be expressed by means of normal relations between node templates, whereas
 455 such modeling would not be possible in **requirements** sections of disjoint node templates.

456 **Example 15 - Referencing an abstract database node template**

```
tosca_definitions_version: tosca_simple_yaml_1_1_0

topology_template:
  description: Template of an application connecting to a database.

  node_templates:
    web_app:
      type: tosca.nodes.WebApplication.MyWebApp
      requirements:
        - host: web_server
        - database_endpoint: db
```

```

web_server:
  type: toska.nodes.WebServer
  requirements:
    - host: server

server:
  type: toska.nodes.Compute
  # details omitted for brevity

db:
  # This node is abstract (no Deployment or Implementation artifacts on
  create)
  # and can be substituted with a topology provided by another template
  # that exports a Database type's capabilities.
  type: toska.nodes.Database
  properties:
    user: my_db_user
    password: secret
    name: my_db_name

```

457 2.10.3 Definition of the database stack in a service template

458 The following sample defines a template for a database including its complete hosting stack, i.e. the
 459 template includes a **database** node template, a template for the database management system (**dbms**)
 460 hosting the database, as well as a computer node **server** on which the DBMS is installed.

461 This service template can be used standalone for deploying just a database and its hosting stack. In the
 462 context of the current use case, though, this template can also substitute the database node template in
 463 the previous snippet and thus fill in the details of how to deploy the database.

464 In order to enable such a substitution, an additional metadata section **substitution_mappings** is added
 465 to the topology template to tell a TOSCA Orchestrator how exactly the topology template will fit into the
 466 context where it gets used. For example, requirements or capabilities of the node that gets substituted by
 467 the topology template have to be mapped to requirements or capabilities of internal node templates for
 468 allow for a proper wiring of the resulting overall graph of node templates.

469 In short, the **substitution_mappings** section provides the following information:

- 470 1. It defines what node templates, i.e. node templates of which type, can be substituted by the
 471 topology template.
- 472 2. It defines how capabilities of the substituted node (or the capabilities defined by the node type of
 473 the substituted node template, respectively) are bound to capabilities of node templates defined
 474 in the topology template.
- 475 3. It defines how requirements of the substituted node (or the requirements defined by the node type
 476 of the substituted node template, respectively) are bound to requirements of node templates
 477 defined in the topology template.

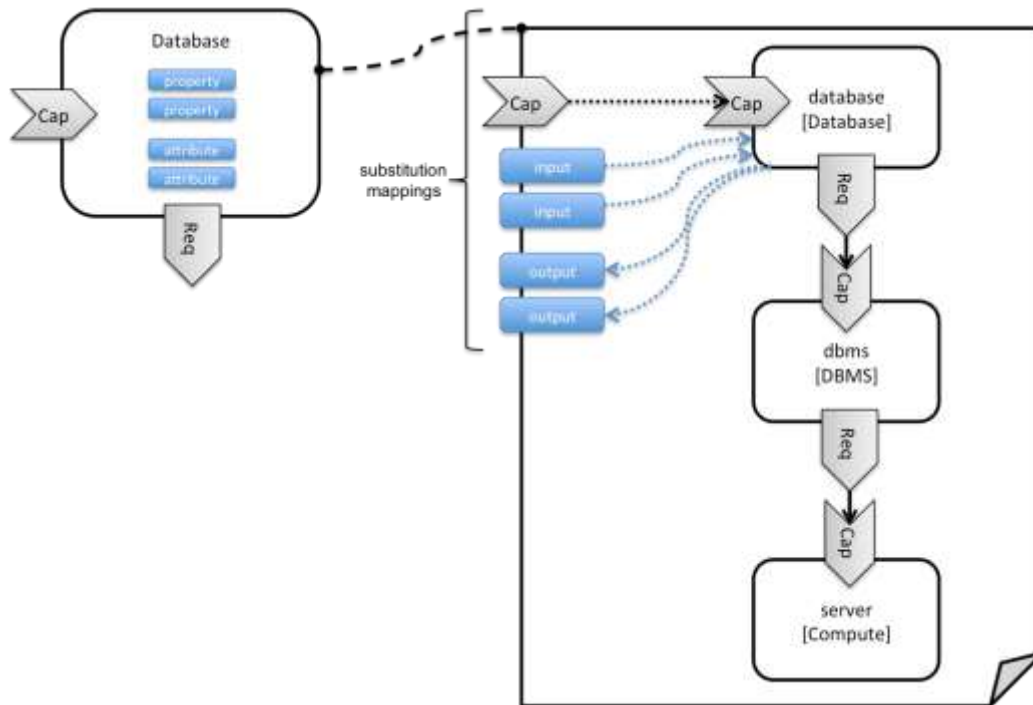


Figure 2: Substitution mappings

478

479

480 The **substitution_mappings** section in the sample below denotes that this topology template can be
 481 used for substituting node templates of type **tosca.nodes.Database**. It further denotes that the
 482 **database_endpoint** capability of the substituted node gets fulfilled by the **database_endpoint**
 483 capability of the **database** node contained in the topology template.

484 **Example 16 - Using substitution mappings to export a database implementation**

```
tosca_definitions_version: tosca_simple_yaml_1_0

topology_template:
  description: Template of a database including its hosting stack.

  inputs:
    db_user:
      type: string
    db_password:
      type: string
    # other inputs omitted for brevity

  substitution_mappings:
    node_type: tosca.nodes.Database
    capabilities:
      database_endpoint: [ database, database_endpoint ]

  node_templates:
    database:
      type: tosca.nodes.Database
      properties:
        user: { get_input: db_user }
        # other properties omitted for brevity
```



```
requirements:
  - host: dbms

dbms:
  type: tosca.nodes.DBMS
  # details omitted for brevity

server:
  type: tosca.nodes.Compute
  # details omitted for brevity
```

485 Note that the **substitution_mappings** section does not define any mappings for requirements of the
486 Database node type, since all requirements are fulfilled by other nodes templates in the current topology
487 template. In cases where a requirement of a substituted node is bound in the top-level service template
488 as well as in the substituting topology template, a TOSCA Orchestrator should raise a validation error.

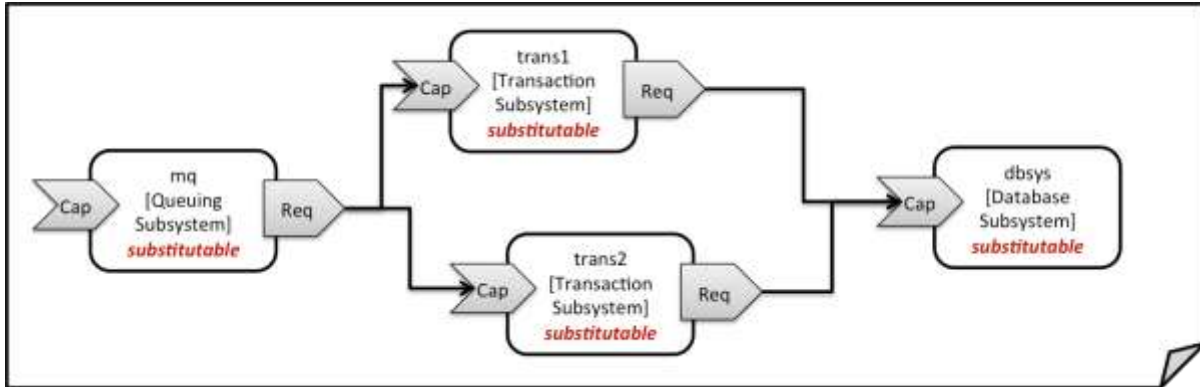
489 Further note that no mappings for properties or attributes of the substituted node are defined. Instead, the
490 inputs and outputs defined by the topology template have to match the properties and attributes of the
491 substituted node. If there are more inputs than the substituted node has properties, default values must
492 be defined for those inputs, since no values can be assigned through properties in a substitution case.

493 2.11 Using node template substitution for chaining subsystems

494 A common use case when providing an end-to-end service is to define a chain of several subsystems that
495 together implement the overall service. Those subsystems are typically defined as separate service
496 templates to (1) keep the complexity of the end-to-end service template at a manageable level and to (2)
497 allow for the re-use of the respective subsystem templates in many different contexts. The type of
498 subsystems may be specific to the targeted workload, application domain, or custom use case. For
499 example, a company or a certain industry might define a subsystem type for company- or industry specific
500 data processing and then use that subsystem type for various end-user services. In addition, there might
501 be generic subsystem types like a database subsystem that are applicable to a wide range of use cases.

502 2.11.1 Defining the overall subsystem chain

503 Figure 3 shows the chaining of three subsystem types – a message queuing subsystem, a transaction
504 processing subsystem, and a databank subsystem – that support, for example, an online booking
505 application. On the front end, this chain provides a capability of receiving messages for handling in the
506 message queuing subsystem. The message queuing subsystem in turn requires a number of receivers,
507 which in the current example are two transaction processing subsystems. The two instances of the
508 transaction processing subsystem might be deployed on two different hosting infrastructures or
509 datacenters for high-availability reasons. The transaction processing subsystems finally require a
510 database subsystem for accessing and storing application specific data. The database subsystem in the
511 backend does not require any further component and is therefore the end of the chain in this example.



512

513

Figure 3: Chaining of subsystems in a service template

514 All of the node templates in the service template shown above are abstract and considered substitutable
 515 where each can be treated as their own subsystem; therefore, when instantiating the overall service, the
 516 orchestrator would realize each substitutable node template using other TOSCA service templates.

517 These service templates would include more nodes and relationships that include the details for each
 518 subsystem. A simplified version of a TOSCA service template for the overall service is given in the
 519 following listing.

520

521 **Example 17 - Declaring a transaction subsystem as a chain of substitutable node templates**

```

tosca_definitions_version: tosca_simple_yaml_1_1_0

topology_template:
  description: Template of online transaction processing service.

  node_templates:
    mq:
      type: example.QueuingSubsystem
      properties:
        # properties omitted for brevity
      capabilities:
        message_queue_endpoint:
          # details omitted for brevity
      requirements:
        - receiver: trans1
        - receiver: trans2

    trans1:
      type: example.TransactionSubsystem
      properties:
        mq_service_ip: { get_attribute: [ mq, service_ip ] }
        receiver_port: 8080
      capabilities:
        message_receiver:
          # details omitted for brevity
      requirements:
        - database_endpoint: dbsys

    trans2:

```

```

type: example.TransactionSubsystem
properties:
  mq_service_ip: { get_attribute: [ mq, service_ip ] }
  receiver_port: 8080
capabilities:
  message_receiver:
    # details omitted for brevity
requirements:
  - database_endpoint: dbsys

dbsys:
type: example.DatabaseSubsystem
properties:
  # properties omitted for brevity
capabilities:
  database_endpoint:
    # details omitted for brevity

```

522

523 As can be seen in the example above, the subsystems are chained to each other by binding requirements
524 of one subsystem node template to other subsystem node templates that provide the respective
525 capabilities. For example, the **receiver** requirement of the message queuing subsystem node template
526 **mq** is bound to transaction processing subsystem node templates **trans1** and **trans2**.

527 Subsystems can be parameterized by providing properties. In the listing above, for example, the IP
528 address of the message queuing server is provided as property **mq_service_ip** to the transaction
529 processing subsystems and the desired port for receiving messages is specified by means of the
530 **receiver_port** property.

531 If attributes of the instantiated subsystems need to be obtained, this would be possible by using the
532 **get_attribute** intrinsic function on the respective subsystem node templates.

533 2.11.2 Defining a subsystem (node) type

534 The types of subsystems that are required for a certain end-to-end service are defined as TOSCA node
535 types as shown in the following example. Node templates of those node types can then be used in the
536 end-to-end service template to define subsystems to be instantiated and chained for establishing the end-
537 to-end service.

538 The realization of the defined node type will be given in the form of a whole separate service template as
539 outlined in the following section.

540

541 Example 18 - Defining a TransactionSubsystem node type

```

tosca_definitions_version: tosca_simple_yaml_1_0

node_types:
  example.TransactionSubsystem:
    properties:
      mq_service_ip:
        type: string
      receiver_port:
        type: integer
    attributes:

```

```

receiver_ip:
  type: string
receiver_port:
  type: integer
capabilities:
  message_receiver: tosca.capabilities.Endpoint
requirements:
  - database_endpoint: tosca.capabilities.Endpoint.Database

```

542

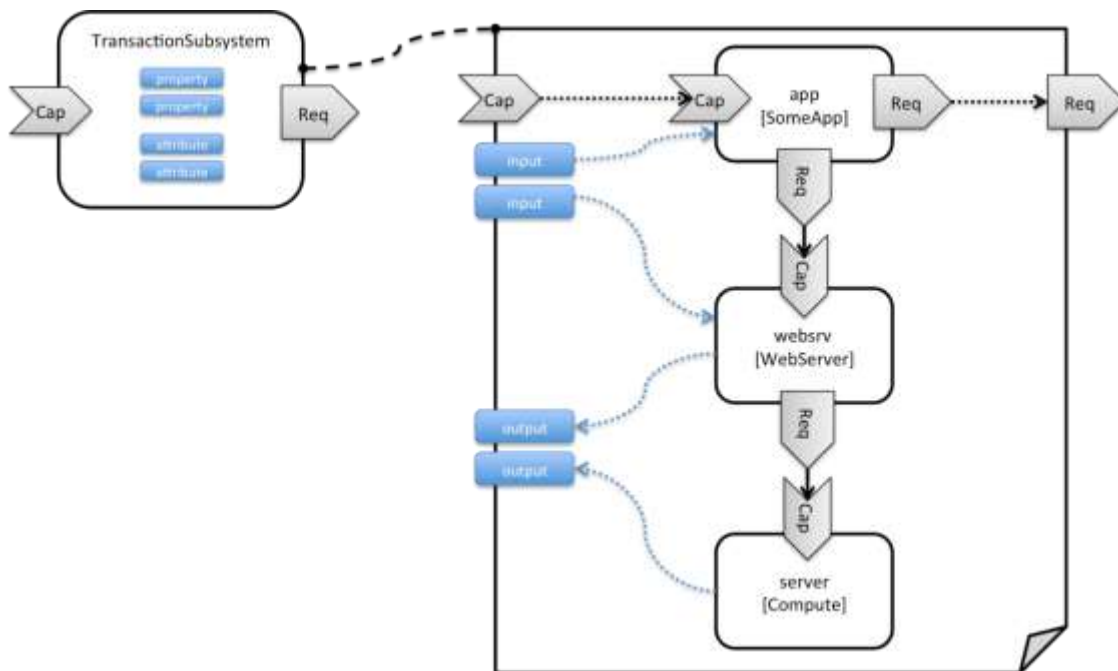
543 Configuration parameters that would be allowed for customizing the instantiation of any subsystem are
544 defined as properties of the node type. In the current example, those are the properties `mq_service_ip`
545 and `receiver_port` that had been used in the end-to-end service template in section 2.11.1.

546 Observable attributes of the resulting subsystem instances are defined as attributes of the node type. In
547 the current case, those are the IP address of the message receiver as well as the actually allocated port
548 of the message receiver endpoint.

549 2.11.3 Defining the details of a subsystem

550 The details of a subsystem, i.e. the software components and their hosting infrastructure, are defined as
551 node templates and relationships in a service template. By means of substitution mappings that have
552 been introduced in section 2.10.2, the service template is annotated to indicate to an orchestrator that it
553 can be used as realization of a node template of certain type, as well as how characteristics of the node
554 type are mapped to internal elements of the service template.

555



556

557 **Figure 4: Defining subsystem details in a service template**

558 Figure 1 illustrates how a transaction processing subsystem as outlined in the previous section could be
559 defined in a service template. In this example, it simply consists of a custom application `app` of type
560 `SomeApp` that is hosted on a web server `websrv`, which in turn is running on a compute node.

561 The application named `app` provides a capability to receive messages, which is bound to the
562 `message_receiver` capability of the substitutable node type. It further requires access to a database, so

563 the application's **database_endpoint** requirement is mapped to the **database_endpoint** requirement of
564 the **TransactionSubsystem** node type.

565 Properties of the **TransactionSubsystem** node type are used to customize the instantiation of a
566 subsystem. Those properties can be mapped to any node template for which the author of the subsystem
567 service template wants to expose configurability. In the current example, the application app and the web
568 server middleware **websrv** get configured through properties of the **TransactionSubsystem** node type.
569 All properties of that node type are defined as **inputs** of the service template. The input parameters in
570 turn get mapped to node templates by means of **get_input** function calls in the respective sections of
571 the service template.

572 Similarly, attributes of the whole subsystem can be obtained from attributes of particular node templates.
573 In the current example, attributes of the web server and the hosting compute node will be exposed as
574 subsystem attributes. All exposed attributes that are defined as attributes of the substitutable
575 **TransactionSubsystem** node type are defined as outputs of the subsystem service template.

576 An outline of the subsystem service template is shown in the listing below. Note that this service template
577 could be used for stand-alone deployment of a transaction processing system as well, i.e. it is not
578 restricted just for use in substitution scenarios. Only the presence of the **substitution_mappings**
579 metadata section in the **topology_template** enables the service template for substitution use cases.

580

581 **Example 19 - Implementation of a TransactionSubsystem node type using substitution mappings**

```
tosca_definitions_version: tosca_simple_yaml_1_0

topology_template:
  description: Template of a database including its hosting stack.

  inputs:
    mq_service_ip:
      type: string
      description: IP address of the message queuing server to receive
messages from
    receiver_port:
      type: string
      description: Port to be used for receiving messages
    # other inputs omitted for brevity

  substitution_mappings:
    node_type: example.TransactionSubsystem
    capabilities:
      message_receiver: [ app, message_receiver ]
    requirements:
      database_endpoint: [ app, database ]

  node_templates:
    app:
      type: example.SomeApp
      properties:
        # properties omitted for brevity
      capabilities:
        message_receiver:
          properties:
```

```

        service_ip: { get_input: mq_service_ip }
        # other properties omitted for brevity
requirements:
  - database:
      # details omitted for brevity
  - host: webserv

webserv:
  type: toska.nodes.WebServer
  properties:
    # properties omitted for brevity
  capabilities:
    data_endpoint:
      properties:
        port_name: { get_input: receiver_port }
        # other properties omitted for brevity
  requirements:
    - host: server

server:
  type: toska.nodes.Compute
  # details omitted for brevity

outputs:
  receiver_ip:
    description: private IP address of the message receiver application
    value: { get_attribute: [ server, private_address ] }
  receiver_port:
    description: Port of the message receiver endpoint
    value: { get_attribute: [ app, app_endpoint, port ] }

```

582 2.12 Grouping node templates

583 In designing applications composed of several interdependent software components (or nodes) it is often
 584 desirable to manage these components as a named group. This can provide an effective way of
 585 associating policies (e.g., scaling, placement, security or other) that orchestration tools can apply to all
 586 the components of group during deployment or during other lifecycle stages.

587 In many realistic scenarios it is desirable to include scaling capabilities into an application to be able to
 588 react on load variations at runtime. The example below shows the definition of a scaling web server stack,
 589 where a variable number of servers with apache installed on them can exist, depending on the load on
 590 the servers.

591 Example 20 - Grouping Node Templates for possible policy application

```

tosca_definitions_version: toska_simple_yaml_1_0

description: Template for a scaling web server.

topology_template:
  inputs:

```

```

# omitted here for brevity

node_templates:
  apache:
    type: tosca.nodes.WebServer.Apache
    properties:
      # Details omitted for brevity
    requirements:
      - host: server

  server:
    type: tosca.nodes.Compute
    # details omitted for brevity

groups:
  webservers_group:
    type: tosca.groups.Root
    members: [ apache, server ]

```

592 The example first of all uses the concept of grouping to express which components (node templates)
593 need to be scaled as a unit – i.e. the compute nodes and the software on-top of each compute node. This
594 is done by defining the **webservers_group** in the **groups** section of the template and by adding both the
595 **apache** node template and the **server** node template as a member to the group.

596 Furthermore, a scaling policy is defined for the group to express that the group as a whole (i.e. pairs of
597 **server** node and the **apache** component installed on top) should scale up or down under certain
598 conditions.

599 In cases where no explicit binding between software components and their hosting compute resources is
600 defined in a template, but only requirements are defined as has been shown in section 2.9, a provider
601 could decide to place software components on the same host if their hosting requirements match, or to
602 place them onto different hosts.

603 It is often desired, though, to influence placement at deployment time to make sure components get
604 collocation or anti-collocated. This can be expressed via grouping and policies as shown in the example
605 below.

606 **Example 21 - Grouping nodes for anti-collocation policy application**

```

tosca_definitions_version: tosca_simple_yaml_1_0

description: Template hosting requirements and placement policy.

topology_template:
  inputs:
    # omitted here for brevity

  node_templates:
    wordpress_server:
      type: tosca.nodes.WebServer
      properties:
        # omitted here for brevity
      requirements:
        - host:
            # Find a Compute node that fulfills these additional filter reqs.
            node_filter:
              capabilities:

```

```

- host:
  properties:
    - mem_size: { greater_or_equal: 512 MB }
    - disk_size: { greater_or_equal: 2 GB }
- os:
  properties:
    - architecture: x86_64
    - type: linux

mysql:
  type: tosca.nodes.DBMS.MySQL
  properties:
    # omitted here for brevity
  requirements:
    - host:
      node: tosca.nodes.Compute
      node_filter:
        capabilities:
          - host:
              properties:
                - disk_size: { greater_or_equal: 1 GB }
            - os:
                properties:
                  - architecture: x86_64
                  - type: linux

groups:
  my_co_location_group:
    type: tosca.groups.Root
    members: [ wordpress_server, mysql ]

policies:
  - my_anti_collocation_policy:
    type: my.policies.anticollocateion
    targets: [ my_co_location_group ]
    # For this example, specific policy definitions are considered
    # domain specific and are not included here

```

607 In the example above, both software components **wordpress_server** and **mysql** have similar hosting
608 requirements. Therefore, a provider could decide to put both on the same server as long as both their
609 respective requirements can be fulfilled. By defining a group of the two components and attaching an anti-
610 collocation policy to the group it can be made sure, though, that both components are put onto different
611 hosts at deployment time.

612 2.13 Using YAML Macros to simplify templates

613 The YAML 1.2 specification allows for defining of [aliases](#), which allow for authoring a block of YAML (or
614 node) once and indicating it is an “anchor” and then referencing it elsewhere in the same document as an
615 “alias”. Effectively, YAML parsers treat this as a “macro” and copy the anchor block’s code to wherever it
616 is referenced. Use of this feature is especially helpful when authoring TOSCA Service Templates where
617 similar definitions and property settings may be repeated multiple times when describing a multi-tier
618 application.

619

620 For example, an application that has a web server and database (i.e., a two-tier application) may be
621 described using two **Compute** nodes (one to host the web server and another to host the database). The

622 author may want both Compute nodes to be instantiated with similar properties such as operating system,
623 distribution, version, etc.

624 To accomplish this, the author would describe the reusable properties using a named anchor in the
625 “**dsl_definitions**” section of the TOSCA Service Template and reference the anchor name as an alias
626 in any **Compute** node templates where these properties may need to be reused. For example:

627 **Example 22 - Using YAML anchors in TOSCA templates**

```
tosca_definitions_version: toska_simple_yaml_1_0

description: >
  TOSCA simple profile that just defines a YAML macro for commonly reused Compute
  properties.

dsl_definitions:
  my_compute_node_props: &my_compute_node_props
    disk_size: 10 GB
    num_cpus: 1
    mem_size: 2 GB

topology_template:
  node_templates:
    my_server:
      type: Compute
      capabilities:
        - host:
            properties: *my_compute_node_props

    my_database:
      type: Compute
      capabilities:
        - host:
            properties: *my_compute_node_props
```

628 **2.14 Passing information as inputs to Nodes and Relationships**

629 It is possible for type and template authors to declare input variables within an **inputs** block on interfaces
630 to nodes or relationships in order to pass along information needed by their operations (scripts). These
631 declarations can be scoped such as to make these variable values available to all operations on a node
632 or relationships interfaces or to individual operations. TOSCA orchestrators will make these values
633 available as environment variables within the execution environments in which the scripts associated with
634 lifecycle operations are run.

635 **2.14.1 Example: declaring input variables for all operations on a single** 636 **interface**

```
node_templates:
  wordpress:
    type: toska.nodes.WebApplication.WordPress
    requirements:
      ...
```



```

- database_endpoint: mysql_database
interfaces:
  Standard:
    inputs:
      wp_db_port: { get_property: [ SELF, database_endpoint, port ] }

```

637 2.14.2 Example: declaring input variables for a single operation

```

node_templates:
  wordpress:
    type: tosca.nodes.WebApplication.WordPress
    requirements:
      ...
      - database_endpoint: mysql_database
    interfaces:
      Standard:
        create: wordpress_install.sh
        configure:
          implementation: wordpress_configure.sh
        inputs:
          wp_db_port: { get_property: [ SELF, database_endpoint, port ] }

```

638 In the case where an input variable name is defined at more than one scope within the same interfaces
639 section of a node or template definition, the lowest (or innermost) scoped declaration would override
640 those declared at higher (or more outer) levels of the definition.

641 2.14.3 Example: setting output variables to an attribute

```

node_templates:
  frontend:
    type: MyTypes.SomeNodeType
    attributes:
      url: { get_operation_output: [ SELF, Standard, create, generated_url ] }
    interfaces:
      Standard:
        create:
          implementation: scripts/frontend/create.sh

```

642
643 In this example, the Standard create operation exposes / exports an environment variable named
644 “**generated_url**” attribute which will be assigned to the WordPress node’s **url** attribute.

645 2.14.4 Example: passing output variables between operations

```

node_templates:
  frontend:
    type: MyTypes.SomeNodeType
    interfaces:
      Standard:
        create:
          implementation: scripts/frontend/create.sh
        configure:
          implementation: scripts/frontend/configure.sh
        inputs:

```

```
    data_dir: { get_operation_output: [ SELF, Standard, create, data_dir ]
  }
```

646 In this example, the **Standard** lifecycle's **create** operation exposes / exports an environment variable
647 named "**data_dir**" which will be passed as an input to the **Standard** lifecycle's **configure** operation.

648 2.15 Topology Template Model versus Instance Model

649 A TOSCA service template contains a **topology template**, which models the components of an
650 application, their relationships and dependencies (a.k.a., a topology model) that get interpreted and
651 instantiated by TOSCA Orchestrators. The actual node and relationship instances that are created
652 represent a set of resources distinct from the template itself, called a **topology instance (model)**. The
653 direction of this specification is to provide access to the instances of these resources for management
654 and operational control by external administrators. This model can also be accessed by an orchestration
655 engine during deployment – i.e. during the actual process of instantiating the template in an incremental
656 fashion. That is, the orchestrator can choose the order of resources to instantiate (i.e., establishing a
657 partial set of node and relationship instances) and have the ability, as they are being created, to access
658 them in order to facilitate instantiating the remaining resources of the complete topology template.

659 2.16 Using attributes implicitly reflected from properties

660 Most entity types in TOSCA (e.g., Node, Relationship, Requirement and Capability Types) have [property](#)
661 [definitions](#), which allow template authors to set the values for as inputs when these entities are
662 instantiated by an orchestrator. These property values are considered to reflect the desired state of the
663 entity by the author. Once instantiated, the actual values for these properties on the realized
664 (instantiated) entity are obtainable via attributes on the entity with the same name as the corresponding
665 property.

666 In other words, TOSCA orchestrators will automatically reflect (i.e., make available) any property defined
667 on an entity making it available as an attribute of the entity with the same name as the property.

668

669 Use of this feature is shown in the example below where a source node named **my_client**, of type
670 **ClientNode**, requires a connection to another node named **my_server** of type **ServerNode**. As you can
671 see, the **ServerNode** type defines a property named **notification_port** which defines a dedicated port
672 number which instances of **my_client** may use to post asynchronous notifications to it during runtime. In
673 this case, the TOSCA Simple Profile assures that the **notification_port** property is implicitly reflected
674 as an attribute in the **my_server** node (also with the name **notification_port**) when its node template
675 is instantiated.

676

677 Example 23 - Properties reflected as attributes

```
tosca_definitions_version: tosca_simple_yaml_1_0

description: >
  TOSCA simple profile that shows how the (notification_port) property is reflected
  as an attribute and can be referenced elsewhere.

node_types:
  ServerNode:
    derived_from: SoftwareComponent
    properties:
      notification_port:
        type: integer
    capabilities:
      # omitted here for brevity
```

```

ClientNode:
  derived_from: SoftwareComponent
  properties:
    # omitted here for brevity
  requirements:
    - server:
        capability: Endpoint
        node: ServerNode
        relationship: ConnectsTo

topology_template:
  node_templates:

    my_server:
      type: ServerNode
      properties:
        notification_port: 8000

    my_client:
      type: ClientNode
      requirements:
        - server:
            node: my_server
            relationship: my_connection

  relationship_templates:
    my_connection:
      type: ConnectsTo
      interfaces:
        Configure:
          inputs:
            targ_notify_port: { get_attribute: [ TARGET, notification_port ]
}

    # other operation definitions omitted here for brevity

```

678

679 Specifically, the above example shows that the **ClientNode** type needs the **notification_port** value
680 anytime a node of **ServerType** is connected to it using the **ConnectsTo** relationship in order to make it
681 available to its **Configure** operations (scripts). It does this by using the **get_attribute** function to
682 retrieve the **notification_port** attribute from the **TARGET** node of the **ConnectsTo** relationship (which is
683 a node of type **ServerNode**) and assigning it to an environment variable named **targ_notify_port**.

684

685 It should be noted that the actual port value of the **notification_port** attribute may or may not be the
686 value **8000** as requested on the property; therefore, any node that is dependent on knowing its actual
687 “runtime” value would use the **get_attribute** function instead of the **get_property** function.

688 3 TOSCA Simple Profile definitions in YAML

689 Except for the examples, this section is **normative** and describes all of the YAML grammar, definitions
690 and block structure for all keys and mappings that are defined for the TOSCA Version 1.0 Simple Profile
691 specification that are needed to describe a TOSCA Service Template (in YAML).

692 3.1 TOSCA Namespace URI and alias

693 The following TOSCA Namespace URI alias and TOSCA Namespace Alias are reserved values which
694 SHALL be used when identifying the TOSCA Simple Profile version 1.0 specification.

Namespace Alias	Namespace URI	Specification Description
tosca_simple_yaml_1_1	http://docs.oasis-open.org/tosca/ns/simple/yaml/1.1	The TOSCA Simple Profile v1.1 (YAML) target namespace and namespace alias.

695 3.1.1 TOSCA Namespace prefix

696 The following TOSCA Namespace prefix is a reserved value and SHALL be used to reference the default
697 TOSCA Namespace URI as declared in TOSCA Service Templates.

Namespace Prefix	Specification Description
tosca	The reserved TOSCA Simple Profile Specification prefix that can be associated with the default TOSCA Namespace URI

698 3.1.2 TOSCA Namespacing in TOSCA Service Templates

699 In the TOSCA Simple Profile, TOSCA Service Templates MUST always have, as the first line of YAML,
700 the keyword "**tosca_definitions_version**" with an associated TOSCA Namespace Alias value. This
701 single line accomplishes the following:

- 702 1. Establishes the TOSCA Simple Profile Specification version whose grammar MUST be used to
703 parse and interpret the contents for the remainder of the TOSCA Service Template.
- 704 2. Establishes the default TOSCA Namespace URI and Namespace Prefix for all types found in the
705 document that are not explicitly namespaced.
- 706 3. Automatically imports (without the use of an explicit import statement) the normative type
707 definitions (e.g., Node, Relationship, Capability, Artifact, etc.) that are associated with the TOSCA
708 Simple Profile Specification the TOSCA Namespace Alias value identifies.
- 709 4. Associates the TOSCA Namespace URI and Namespace Prefix to the automatically imported
710 TOSCA type definitions.

711 3.1.3 Rules to avoid namespace collisions

712 TOSCA Simple Profiles allows template authors to declare their own types and templates and assign
713 them simple names with no apparent namespaces. Since TOSCA Service Templates can import other
714 service templates to introduce new types and topologies of templates that can be used to provide
715 concrete implementations (or substitute) for abstract nodes. Rules are needed so that TOSCA
716 Orchestrators know how to avoid collisions and apply their own namespaces when import and nesting
717 occur.

718 3.1.3.1 Additional Requirements

- 719
- 720 • Since TOSCA Service Templates can import (or substitute in) other Service Templates, TOSCA
721 Orchestrators and tooling will encounter the “`tosca_definitions_version`” statement for each
722 imported template. In these cases, the following additional requirements apply:
 - 723 ○ Imported type definitions with the same Namespace URI, local name and version SHALL
724 be equivalent.
 - 725 ○ If different values of the “`tosca_definitions_version`” are encountered, their
726 corresponding type definitions MUST be uniquely identifiable using their corresponding
727 Namespace URI using a different Namespace prefix.
 - 728 • Duplicate local names (i.e., within the same Service Template SHALL be considered an error.
729 These include, but are not limited to duplicate names found for the following definitions:
 - 730 ○ Repositories (repositories)
 - 731 ○ Data Types (data_types)
 - 732 ○ Node Types (node_types)
 - 733 ○ Relationship Types (relationship_types)
 - 734 ○ Capability Types (capability_types)
 - 735 ○ Artifact Types (artifact_types)
 - 736 ○ Interface Types (interface_types)
 - 737 • Duplicate Template names within a Service Template’s Topology Template SHALL be considered
738 an error. These include, but are not limited to duplicate names found for the following template
739 types:
 - 740 ○ Node Templates (node_templates)
 - 741 ○ Relationship Templates (relationship_templates)
 - 742 ○ Inputs (inputs)
 - 743 ○ Outputs (outputs)
 - 744 • Duplicate names for the following keynames within Types or Templates SHALL be considered an
745 error. These include, but are not limited to duplicate names found for the following keynames:
 - 746 ○ Properties (properties)
 - 747 ○ Attributes (attributes)
 - 748 ○ Artifacts (artifacts)
 - 749 ○ Requirements (requirements)
 - 750 ○ Capabilities (capabilities)
 - 751 ○ Interfaces (interfaces)
 - 752 ○ Policies (policies)
 - Groups (groups)

753 3.2 Parameter and property types

754 This clause describes the primitive types that are used for declaring normative properties, parameters
755 and grammar elements throughout this specification.

756 3.2.1 Referenced YAML Types

757 Many of the types we use in this profile are built-in types from the [YAML 1.2 specification](#) (i.e., those
758 identified by the “tag:yaml.org,2002” version tag) [[YAML-1.2](#)].

759 The following table declares the valid YAML type URIs and aliases that SHALL be used when possible
760 when defining parameters or properties within TOSCA Service Templates using this specification:

Valid aliases	Type URI
string	tag:yaml.org,2002:str (default)

integer	tag:yaml.org,2002:int
float	tag:yaml.org,2002:float
boolean	tag:yaml.org,2002:bool (i.e., a value either 'true' or 'false')
timestamp	tag:yaml.org,2002:timestamp [YAML-TS-1.1]
null	tag:yaml.org,2002:null

761 3.2.1.1 Notes

- 762
- The “string” type is the default type when not specified on a parameter or property declaration.
 - While YAML supports further type aliases, such as “str” for “string”, the TOSCA Simple Profile specification promotes the fully expressed alias name for clarity.
- 763
- 764

765 3.2.2 TOSCA version

766 TOSCA supports the concept of “reuse” of type definitions, as well as template definitions which could be
767 version and change over time. It is important to provide a reliable, normative means to represent a
768 version string which enables the comparison and management of types and templates over time.
769 Therefore, the TOSCA TC intends to provide a normative version type (string) for this purpose in future
770 Working Drafts of this specification.

Shorthand Name	version
Type Qualified Name	tosca:version

771 3.2.2.1 Grammar

772 TOSCA version strings have the following grammar:

```
<major_version>.<minor_version>[.<fix_version>[.<qualifier>[-<build_version>] ] ]
```

773 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 774
- **major_version**: is a required integer value greater than or equal to 0 (zero)
 - **minor_version**: is a required integer value greater than or equal to 0 (zero).
 - **fix_version**: is an optional integer value greater than or equal to 0 (zero).
 - **qualifier**: is an optional string that indicates a named, pre-release version of the associated code that has been derived from the version of the code identified by the combination **major_version**, **minor_version** and **fix_version** numbers.
 - **build_version**: is an optional integer value greater than or equal to 0 (zero) that can be used to further qualify different build versions of the code that has the same **qualifier_string**.
- 775
- 776
- 777
- 778
- 779
- 780
- 781

782 3.2.2.2 Version Comparison

- 783
- When comparing TOSCA versions, all component versions (i.e., *major*, *minor* and *fix*) are compared in sequence from left to right.
 - TOSCA versions that include the optional qualifier are considered older than those without a qualifier.
 - TOSCA versions with the same major, minor, and fix versions and have the same qualifier string, but with different build versions can be compared based upon the build version.
 - Qualifier strings are considered domain-specific. Therefore, this specification makes no recommendation on how to compare TOSCA versions with the same major, minor and fix
- 784
- 785
- 786
- 787
- 788
- 789
- 790

791 versions, but with different qualifiers strings and simply considers them different named branches
792 derived from the same code.

793 3.2.2.3 Examples

794 Examples of valid TOSCA version strings:

```
# basic version strings
6.1
2.0.1

# version string with optional qualifier
3.1.0.beta

# version string with optional qualifier and build version
1.0.0.alpha-10
```

795 3.2.2.4 Notes

796 • [\[Maven-Version\]](#) The TOSCA version type is compatible with the Apache Maven versioning
797 policy.

798 3.2.2.5 Additional Requirements

- 799 • A version value of zero (i.e., '0', '0.0', or '0.0.0') SHALL indicate there no version provided.
800 • A version value of zero used with any qualifiers SHALL NOT be valid.

801 3.2.3 TOSCA range type

802 The range type can be used to define numeric ranges with a lower and upper boundary. For example, this
803 allows for specifying a range of ports to be opened in a firewall.

Shorthand Name	range
Type Qualified Name	tosca:range

804 3.2.3.1 Grammar

805 TOSCA range values have the following grammar:

```
[<lower_bound>, <upper_bound>]
```

806 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 807 • **lower_bound**: is a required integer value that denotes the lower boundary of the range.
808 • **upper_bound**: is a required integer value that denotes the upper boundary of the range. This
809 value **MUST** be greater than **lower_bound**.

810 3.2.3.2 Keywords

811 The following Keywords may be used in the TOSCA range type:

Keyword	Applicable Types	Description
UNBOUNDED	scalar	Used to represent an unbounded upper bounds (positive) value in a set for a scalar type.

812 3.2.3.3 Examples

813 Example of a node template property with a range value:

```
# numeric range between 1 and 100
a_range_property: [ 1, 100 ]

# a property that has allows any number 0 or greater
num_connections: [ 0, UNBOUNDED ]
```

814

815 3.2.4 TOSCA list type

816 The list type allows for specifying multiple values for a parameter of property. For example, if an
817 application allows for being configured to listen on multiple ports, a list of ports could be configured using
818 the list data type.

819 Note that entries in a list for one property or parameter must be of the same type. The type (for simple
820 entries) or schema (for complex entries) is defined by the **entry_schema** attribute of the respective
821 [property definition](#), [attribute definitions](#), or input or output [parameter definitions](#).

Shorthand Name	list
Type Qualified Name	tosca:list

822 3.2.4.1 Grammar

823 TOSCA lists are essentially normal YAML lists with the following grammars:

824 3.2.4.1.1 Square bracket notation

```
[ <list_entry_1>, <list_entry_2>, ... ]
```

825 3.2.4.1.2 Bulleted (sequenced) list notation

```
- <list_entry_1>
- ...
- <list_entry_n>
```

826 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 827
- **<list_entry_*>**: represents one entry of the list.

828 3.2.4.2 Declaration Examples

829 3.2.4.2.1 List declaration using a simple type

830 The following example shows a list declaration with an entry schema based upon a simple integer type
831 (which has additional constraints):

```
<some_entity>:
  ...
  properties:
    listen_ports:
      type: list
      entry_schema:
        description: listen port entry (simple integer type)
        type: integer
```



```
constraints:
  - max_length: 128
```

832 3.2.4.2.2 List declaration using a complex type

833 The following example shows a list declaration with an entry schema based upon a complex type:

```
<some_entity>:
  ...
  properties:
    products:
      type: list
      entry_schema:
        description: Product information entry (complex type) defined elsewhere
        type: ProductInfo
```

834 3.2.4.3 Definition Examples

835 These examples show two notation options for defining lists:

- 836 • A single-line option which is useful for only short lists with simple entries.
- 837 • A multi-line option where each list entry is on a separate line; this option is typically useful or
838 more readable if there is a large number of entries, or if the entries are complex.

839 3.2.4.3.1 Square bracket notation

```
listen_ports: [ 80, 8080 ]
```

840 3.2.4.3.2 Bulleted list notation

```
listen_ports:
  - 80
  - 8080
```

841 3.2.5 TOSCA map type

842 The map type allows for specifying multiple values for a parameter of property as a map. In contrast to
843 the list type, where each entry can only be addressed by its index in the list, entries in a map are named
844 elements that can be addressed by their keys.

845 Note that entries in a map for one property or parameter must be of the same type. The type (for simple
846 entries) or schema (for complex entries) is defined by the **entry_schema** attribute of the respective
847 [property definition](#), [attribute definition](#), or input or output [parameter definition](#).

Shorthand Name	map
Type Qualified Name	tosca:map

848 3.2.5.1 Grammar

849 TOSCA maps are normal YAML dictionaries with following grammar:

850 3.2.5.1.1 Single-line grammar

```
{ <entry_key_1>: <entry_value_1>, ..., <entry_key_n>: <entry_value_n> }
...
```

```
<entry_key_n>: <entry_value_n>
```

851 3.2.5.1.2 Multi-line grammar

```
<entry_key_1>: <entry_value_1>  
...  
<entry_key_n>: <entry_value_n>
```

852 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 853 • **entry_key_***: is the required key for an entry in the map
- 854 • **entry_value_***: is the value of the respective entry in the map

855 3.2.5.2 Declaration Examples

856 3.2.5.2.1 Map declaration using a simple type

857 The following example shows a map with an entry schema definition based upon an existing string type
858 (which has additional constraints):

```
<some_entity>:  
...  
properties:  
  emails:  
    type: map  
    entry_schema:  
      description: basic email address  
      type: string  
      constraints:  
        - max_length: 128
```

859 3.2.5.2.2 Map declaration using a complex type

860 The following example shows a map with an entry schema definition for contact information:

```
<some_entity>:  
...  
properties:  
  contacts:  
    type: map  
    entry_schema:  
      description: simple contact information  
      type: ContactInfo
```

861 3.2.5.3 Definition Examples

862 These examples show two notation options for defining maps:

- 863 • A single-line option which is useful for only short maps with simple entries.
- 864 • A multi-line option where each map entry is on a separate line; this option is typically useful or
865 more readable if there is a large number of entries, or if the entries are complex.

866 3.2.5.3.1 Single-line notation

```
# notation option for shorter maps  
user_name_to_id_map: { user1: 1001, user2: 1002 }
```

867 **3.2.5.3.2 Multi-line notation**

```
# notation for longer maps
user_name_to_id_map:
  user1: 1001
  user2: 1002
```

868 **3.2.6 TOSCA scalar-unit type**

869 The scalar-unit type can be used to define scalar values along with a unit from the list of recognized units
870 provided below.

871 **3.2.6.1 Grammar**

872 TOSCA scalar-unit typed values have the following grammar:

```
<scalar> <unit>
```

873 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 874 • **scalar**: is a required scalar value.
- 875 • **unit**: is a required unit value. The unit value **MUST** be type-compatible with the scalar.

876 **3.2.6.2 Additional requirements**

- 877 • **Whitespace**: any number of spaces (including zero or none) **SHALL** be allowed between the
878 **scalar** value and the **unit** value.
- 879 • It **SHALL** be considered an error if either the scalar or unit portion is missing on a property or
880 attribute declaration derived from any scalar-unit type.
- 881 • When performing constraint clause evaluation on values of the scalar-unit type, both the scalar
882 value portion and unit value portion **SHALL** be compared together (i.e., both are treated as a
883 single value). For example, if we have a property called **storage_size**. which is of type scalar-
884 unit, a valid range constraint would appear as follows:
885 o **storage_size**: **in_range** [4 GB, 20 GB]

886 where **storage_size**'s range would be evaluated using both the numeric and unit values
887 (combined together), in this case '4 GB' and '20 GB'.

888 **3.2.6.3 Concrete Types**

Shorthand Names	scalar-unit.size, scalar-unit.time
Type Qualified Names	tosca:scalar-unit.size, toasca:scalar-unit.time

- 889
- 890 The scalar-unit type grammar is abstract and has two recognized concrete types in TOSCA:
- 891 • **scalar-unit.size** – used to define properties that have scalar values measured in size units.
 - 892 • **scalar-unit.time** – used to define properties that have scalar values measured in size units.
 - 893 • **scalar-unit.frequency** – used to define properties that have scalar values measured in units per
894 second.

895 These types and their allowed unit values are defined below.

896 **3.2.6.4 scalar-unit.size**

897 **3.2.6.4.1 Recognized Units**

Unit	Usage	Description
B	size	byte
kB	size	kilobyte (1000 bytes)
KiB	size	kibibytes (1024 bytes)
MB	size	megabyte (1000000 bytes)
MiB	size	mebibyte (1048576 bytes)
GB	size	gigabyte (1000000000 bytes)
GiB	size	gibibytes (1073741824 bytes)
TB	size	terabyte (1000000000000 bytes)
TiB	size	tebibyte (1099511627776 bytes)

898 **3.2.6.4.2 Examples**

```
# Storage size in Gigabytes
properties:
  storage_size: 10 GB
```

899 **3.2.6.4.3 Notes**

- 900 • The unit values recognized by TOSCA Simple Profile for size-type units are based upon a
901 subset of those defined by GNU at
902 http://www.gnu.org/software/parted/manual/html_node/unit.html, which is a non-normative
903 reference to this specification.
- 904 • TOSCA treats these unit values as case-insensitive (e.g., a value of 'kB', 'KB' or 'kb' would be
905 equivalent), but it is considered best practice to use the case of these units as prescribed by
906 GNU.
- 907 • Some Cloud providers may not support byte-level granularity for storage size allocations. In
908 those cases, these values could be treated as desired sizes and actual allocations would be
909 based upon individual provider capabilities.

910 **3.2.6.5 scalar-unit.time**

911 **3.2.6.5.1 Recognized Units**

Unit	Usage	Description
d	time	days
h	time	hours
m	time	minutes

Unit	Usage	Description
s	time	seconds
ms	time	milliseconds
us	time	microseconds
ns	time	nanoseconds

912 3.2.6.5.2 Examples

```
# Response time in milliseconds
properties:
  response_time: 10 ms
```

913 3.2.6.5.3 Notes

- 914 • The unit values recognized by TOSCA Simple Profile for time-type units are based upon a subset
- 915 of those defined by International System of Units whose recognized abbreviations are defined
- 916 within the following reference:
- 917 ○ <http://www.ewh.ieee.org/soc/ias/pub-dept/abbreviation.pdf>
- 918 ○ This document is a non-normative reference to this specification and intended for publications
- 919 or grammars enabled for Latin characters which are not accessible in typical programming
- 920 languages

921 3.2.6.6 scalar-unit.frequency

922 3.2.6.6.1 Recognized Units

Unit	Usage	Description
Hz	frequency	Hertz, or Hz. equals one cycle per second.
kHz	frequency	Kilohertz, or kHz, equals to 1,000 Hertz
MHz	frequency	Megahertz, or MHz, equals to 1,000,000 Hertz or 1,000 kHz
GHz	frequency	Gigahertz, or GHz, equals to 1,000,000,000 Hertz, or 1,000,000 kHz, or 1,000 MHz.

923 3.2.6.6.2 Examples

```
# Processor raw clock rate
properties:
  clock_rate: 2.4 GHz
```

924 3.2.6.6.3 Notes

- 925 • The value for Hertz (Hz) is the International Standard Unit (ISU) as described by the Bureau
- 926 International des Poids et Mesures (BIPM) in the “*SI Brochure: The International System of Units*
- 927 (*SI*) [8th edition, 2006; updated in 2014]”, <http://www.bipm.org/en/publications/si-brochure/>

928 **3.3 Normative values**

929 **3.3.1 Node States**

930 As components (i.e., nodes) of TOSCA applications are deployed, instantiated and orchestrated over
931 their lifecycle using normative lifecycle operations (see section 5.8 for normative lifecycle definitions) it is
932 important define normative values for communicating the states of these components normatively
933 between orchestration and workflow engines and any managers of these applications.

934 The following table provides the list of recognized node states for TOSCA Simple Profile that would be set
935 by the orchestrator to describe a node instance's state:

Node State		
Value	Transitional	Description
initial	no	Node is not yet created. Node only exists as a template definition.
creating	yes	Node is transitioning from initial state to created state.
created	no	Node software has been installed.
configuring	yes	Node is transitioning from created state to configured state.
configured	no	Node has been configured prior to being started.
starting	yes	Node is transitioning from configured state to started state.
started	no	Node is started.
stopping	yes	Node is transitioning from its current state to a configured state.
deleting	yes	Node is transitioning from its current state to one where it is deleted and its state is no longer tracked by the instance model.
error	no	Node is in an error state.

936 **3.3.2 Relationship States**

937 Similar to the Node States described in the previous section, Relationships have state relative to their
938 (normative) lifecycle operations.

939 The following table provides the list of recognized relationship states for TOSCA Simple Profile that would
940 be set by the orchestrator to describe a node instance's state:

Node State		
Value	Transitional	Description
initial	no	Relationship is not yet created. Relationship only exists as a template definition.

941 **3.3.2.1 Notes**

- 942
- Additional states may be defined in future versions of the TOSCA Simple Profile in YAML
943 specification.

944 **3.3.3 Directives**

945 There are currently no directive values defined for this version of the TOSCA Simple Profile.

946 **3.3.4 Network Name aliases**

947 The following are recognized values that may be used as aliases to reference types of networks within an
948 application model without knowing their actual name (or identifier) which may be assigned by the
949 underlying Cloud platform at runtime.

Alias value	Description
PRIVATE	An alias used to reference the first private network within a property or attribute of a Node or Capability which would be assigned to them by the underlying platform at runtime. A private network contains IP addresses and ports typically used to listen for incoming traffic to an application or service from the Intranet and not accessible to the public internet.
PUBLIC	An alias used to reference the first public network within a property or attribute of a Node or Capability which would be assigned to them by the underlying platform at runtime. A public network contains IP addresses and ports typically used to listen for incoming traffic to an application or service from the Internet.

950 **3.3.4.1 Usage**

951 These aliases would be used in the **tosca.capabilities.Endpoint** Capability type (and types derived
952 from it) within the **network_name** field for template authors to use to indicate the type of network the
953 Endpoint is supposed to be assigned an IP address from.

954 **3.4 TOSCA Metamodel**

955 This section defines all modelable entities that comprise the TOSCA Version 1.0 Simple Profile
956 specification along with their keynames, grammar and requirements.

957 **3.4.1 Required Keynames**

958 The TOSCA metamodel includes complex types (e.g., Node Types, Relationship Types, Capability Types,
959 Data Types, etc.) each of which include their own list of reserved keynames that are sometimes marked
960 as **required**. These types may be used to derive other types. These derived types (e.g., child types) do
961 not have to provide required keynames as long as they have been specified in the type they have been
962 derived from (i.e., their parent type).

963 **3.5 Reusable modeling definitions**

964 **3.5.1 Description definition**

965 This optional element provides a means include single or multiline descriptions within a TOSCA Simple
966 Profile template as a scalar string value.

967 **3.5.1.1 Keyname**

968 The following keyname is used to provide a description within the TOSCA Simple Profile specification:

```
description
```

969 **3.5.1.2 Grammar**

970 Description definitions have the following grammar:

```
description: <string>
```

971 **3.5.1.3 Examples**

972 Simple descriptions are treated as a single literal that includes the entire contents of the line that
973 immediately follows the **description** key:

```
description: This is an example of a single line description (no folding).
```

974 The YAML “folded” style may also be used for multi-line descriptions which “folds” line breaks as space
975 characters.

```
description: >  
This is an example of a multi-line description using YAML. It permits for line  
breaks for easier readability...  
  
if needed. However, (multiple) line breaks are folded into a single space  
character when processed into a single string value.
```

976 **3.5.1.4 Notes**

- 977 • Use of “folded” style is discouraged for the YAML string type apart from when used with the
978 **description** keyname.

979 **3.5.2 Constraint clause**

980 A constraint clause defines an operation along with one or more compatible values that can be used to
981 define a constraint on a property or parameter’s allowed values when it is defined in a TOSCA Service
982 Template or one of its entities.

983 **3.5.2.1 Operator keynames**

984 The following is the list of recognized operators (keynames) when defining constraint clauses:

Operator	Type	Value Type	Description
equal	scalar	any	Constrains a property or parameter to a value equal to ('=') the value declared.
greater_than	scalar	comparable	Constrains a property or parameter to a value greater than ('>') the value declared.
greater_or_equal	scalar	comparable	Constrains a property or parameter to a value greater than or equal to ('>=') the value declared.
less_than	scalar	comparable	Constrains a property or parameter to a value less than ('<') the value declared.
less_or_equal	scalar	comparable	Constrains a property or parameter to a value less than or equal to ('<=') the value declared.
in_range	dual scalar	comparable, range	Constrains a property or parameter to a value in range of (inclusive) the two values declared. Note: subclasses or templates of types that declare a property with the in_range constraint MAY only further restrict the range specified by the parent type.

Operator	Type	Value Type	Description
valid_values	list	any	Constrains a property or parameter to a value that is in the list of declared values.
length	scalar	string, list, map	Constrains the property or parameter to a value of a given length.
min_length	scalar	string, list, map	Constrains the property or parameter to a value to a minimum length.
max_length	scalar	string, list, map	Constrains the property or parameter to a value to a maximum length.
pattern	regex	string	Constrains the property or parameter to a value that is allowed by the provided regular expression. Note: Future drafts of this specification will detail the use of regular expressions and reference an appropriate standardized grammar.

985 3.5.2.1.1 Comparable value types

986 In the Value Type column above, an entry of “comparable” includes [integer](#), [float](#), [timestamp](#), [string](#),
987 [version](#), and [scalar-unit](#) types while an entry of “any” refers to any type allowed in the TOSCA simple
988 profile in YAML.

989 3.5.2.2 Additional Requirements

- 990 • If no operator is present for a simple scalar-value on a constraint clause, it **SHALL** be interpreted
991 as being equivalent to having the “**equal**” operator provided; however, the “**equal**” operator may
992 be used for clarity when expressing a constraint clause.
- 993 • The “**length**” operator **SHALL** be interpreted mean “size” for set types (i.e., list, map, etc.).
- 994 • Values provided by the operands (i.e., values and scalar values) **SHALL** be type-compatible with
995 their associated operations.
- 996 • Future drafts of this specification will detail the use of regular expressions and reference an
997 appropriate standardized grammar.

998 3.5.2.3 Grammar

999 Constraint clauses have one of the following grammars:

```
# Scalar grammar
<operator>: <scalar_value>

# Dual scalar grammar
<operator>: [ <scalar_value_1>, <scalar_value_2> ]

# List grammar
<operator> [ <value_1>, <value_2>, ..., <value_n> ]

# Regular expression (regex) grammar
pattern: <regular_expression_value>
```

1000 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1001 • **operator**: represents a required operator from the specified list shown above (see section
1002 3.5.2.1 “Operator keynames”).

- 1003 • **scalar_value**, **scalar_value_***: represents a required scalar (or atomic quantity) that can
- 1004 hold only one value at a time. This will be a value of a primitive type, such as an integer or string
- 1005 that is allowed by this specification.
- 1006 • **value_***: represents a required value of the operator that is not limited to scalars.
- 1007 • **regular_expression_value**: represents a regular expression (string) value.

1008 3.5.2.4 Examples

1009 Constraint clauses used on parameter or property definitions:

```
# equal
equal: 2

# greater_than
greater_than: 1

# greater_or_equal
greater_or_equal: 2

# less_than
less_than: 5

# less_or_equal
less_or_equal: 4

# in_range
in_range: [ 1, 4 ]

# valid_values
valid_values: [ 1, 2, 4 ]
# specific length (in characters)
length: 32

# min_length (in characters)
min_length: 8

# max_length (in characters)
max_length: 64
```

1010 3.5.3 Property Filter definition

1011 A property filter definition defines criteria, using constraint clauses, for selection of a TOSCA entity based
1012 upon its property values.

1013 3.5.3.1 Grammar

1014 Property filter definitions have one of the following grammars:

1015 3.5.3.1.1 Short notation:

1016 The following single-line grammar may be used when only a single constraint is needed on a property:

```
<property_name>: <property constraint clause>
```

1017 3.5.3.1.2 Extended notation:

1018 The following multi-line grammar may be used when multiple constraints are needed on a property:

```

<property_name>:
- <property_constraint_clause_1>
- ...
- <property_constraint_clause_n>

```

1019 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1020 • **property_name**: represents the name of property that would be used to select a property
- 1021 definition with the same name (**property_name**) on a TOSCA entity (e.g., a Node Type, Node
- 1022 Template, Capability Type, etc.).
- 1023 • **property_constraint_clause_***: represents constraint clause(s) that would be used to filter
- 1024 entities based upon the named property's value(s).

1025 3.5.3.2 Additional Requirements

- 1026 • Property constraint clauses must be type compatible with the property definitions (of the same
- 1027 name) as defined on the target TOSCA entity that the clause would be applied against.

1028 3.5.4 Node Filter definition

1029 A node filter definition defines criteria for selection of a TOSCA Node Template based upon the
 1030 template's property values, capabilities and capability properties.

1031 3.5.4.1 Keynames

1032 The following is the list of recognized keynames for a TOSCA node filter definition:

Keyname	Required	Type	Description
properties	no	list of property filter definition	An optional sequenced list of property filters that would be used to select (filter) matching TOSCA entities (e.g., Node Template, Node Type, Capability Types, etc.) based upon their property definitions' values.
capabilities	no	list of capability names or capability type names	An optional sequenced list of capability names or types that would be used to select (filter) matching TOSCA entities based upon their existence.

1033 3.5.4.2 Additional filtering on named Capability properties

1034 Capabilities used as filters often have their own sets of properties which also can be used to construct a
 1035 filter.

Keyname	Required	Type	Description
<capability name_or_type> name>: properties	no	list of property filter definitions	An optional sequenced list of property filters that would be used to select (filter) matching TOSCA entities (e.g., Node Template, Node Type, Capability Types, etc.) based upon their capabilities' property definitions' values.

1036 3.5.4.3 Grammar

1037 Node filter definitions have following grammar:

```

<filter_name>:
  properties:
  - <property_filter_def_1>

```

```

- ...
- <property_filter_def_n>
capabilities:
- <capability_name_or_type_1>:
  properties:
  - <cap_1_property_filter_def_1>
  - ...
  - <cap_m_property_filter_def_n>
- ...
- <capability_name_or_type_n>:
  properties:
  - <cap_1_property_filter_def_1>
  - ...
  - <cap_m_property_filter_def_n>

```

1038 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1039 • **property_filter_def_***: represents a property filter definition that would be used to select
- 1040 (filter) matching TOSCA entities (e.g., Node Template, Node Type, Capability Types, etc.) based
- 1041 upon their property definitions' values.
- 1042 • **capability_name_or_type_***: represents the type or name of a capability that would be used
- 1043 to select (filter) matching TOSCA entities based upon their existence.
- 1044 • **cap_*_property_def_***: represents a property filter definition that would be used to select
- 1045 (filter) matching TOSCA entities (e.g., Node Template, Node Type, Capability Types, etc.) based
- 1046 upon their capabilities' property definitions' values.

1047 3.5.4.4 Additional requirements

- 1048 • TOSCA orchestrators **SHALL** search for matching capabilities listed on a target filter by assuming
- 1049 the capability name is first a symbolic name and secondly it is a type name (in order to avoid
- 1050 namespace collisions).

1051 3.5.4.5 Example

1052 The following example is a filter that would be used to select a TOSCA [Compute](#) node based upon the

1053 values of its defined capabilities. Specifically, this filter would select Compute nodes that supported a

1054 specific range of CPUs (i.e., `num_cpus` value between 1 and 4) and memory size (i.e., `mem_size` of 2 or

1055 greater) from its declared "host" capability.

1056

```

my_node_template:
  # other details omitted for brevity
  requirements:
  - host:
    node_filter:
      capabilities:
        # My "host" Compute node needs these properties:
        - host:
          properties:
            - num_cpus: { in_range: [ 1, 4 ] }
            - mem_size: { greater_or_equal: 512 MB }

```

1057 3.5.5 Repository definition

1058 A repository definition defines a named external repository which contains deployment and

1059 implementation artifacts that are referenced within the TOSCA Service Template.

1060 3.5.5.1 Keynames

1061 The following is the list of recognized keynames for a TOSCA repository definition:

Keyname	Required	Type	Constraints	Description
description	no	description	None	The optional description for the repository.
url	yes	string	None	The required URL or network address used to access the repository.
credential	no	Credential	None	The optional Credential used to authorize access to the repository.

1062 3.5.5.2 Grammar

1063 Repository definitions have one the following grammars:

1064 3.5.5.2.1 Single-line grammar (no credential):

```
<repository\_name>: <repository_address>
```

1065 3.5.5.2.2 Multi-line grammar

```
<repository\_name>:  
  description: <repository\_description>  
  url: <repository\_address>  
  credential: <authorization\_credential>
```

1066 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1067 • **repository_name**: represents the required symbolic name of the repository as a [string](#).
- 1068 • **repository_description**: contains an optional description of the repository.
- 1069 • **repository_address**: represents the required URL of the repository as a string.
- 1070 • **authorization_credential**: represents the optional credentials (e.g., user ID and password)
- 1071 used to authorize access to the repository.

1072 3.5.5.3 Example

1073 The following represents a repository definition:

```
repositories:  
  my_code_repo:  
    description: My project's code repository in GitHub  
    url: https://github.com/my-project/
```

1074 3.5.6 Artifact definition

1075 An artifact definition defines a named, typed file that can be associated with Node Type or Node
1076 Template and used by orchestration engine to facilitate deployment and implementation of interface
1077 operations.

1078 3.5.6.1 Keynames

1079 The following is the list of recognized keynames for a TOSCA artifact definition when using the extended
1080 notation:

Keyname	Required	Type	Description
type	yes	string	The required artifact type for the artifact definition.

Keyname	Required	Type	Description
file	yes	string	The required URI string (relative or absolute) which can be used to locate the artifact's file.
repository	no	string	The optional name of the repository definition which contains the location of the external repository that contains the artifact. The artifact is expected to be referenceable by its file URI within the repository.
description	no	description	The optional description for the artifact definition.
deploy_path	no	string	The file path the associated file would be deployed into within the target node's container.

1081 3.5.6.2 Grammar

1082 Artifact definitions have one of the following grammars:

1083 3.5.6.2.1 Short notation

1084 The following single-line grammar may be used when the artifact's type and mime type can be inferred
1085 from the file URI:

```
<artifact name>: <artifact file URI>
```

1086 3.5.6.2.2 Extended notation:

1087 The following multi-line grammar may be used when the artifact's definition's type and mime type need to
1088 be explicitly declared:

```
<artifact name>:
  description: <artifact description>
  type: <artifact type name>
  file: <artifact file URI>
  repository: <artifact repository name>
  deploy_path: <file deployment path>
```

1089 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1090 • **artifact_name**: represents the required symbolic name of the artifact as a **string**.
- 1091 • **artifact_description**: represents the optional **description** for the artifact.
- 1092 • **artifact_type_name**: represents the required **artifact type** the artifact definition is based upon.
- 1093 • **artifact_file_URI**: represents the required **URI string (relative or absolute) which can be**
1094 **used to locate the artifact's file.**
- 1095 • **artifact_repository_name**: represents the optional name of the **repository definition** to use to
1096 retrieve the associated artifact (file) from.
- 1097 • **file_deployment_path**: represents the optional path the **artifact_file_URI** would be
1098 copied into within the target node's container.

1099 3.5.6.3 Example

1100 The following represents an artifact definition:

```
my_file_artifact: ../my_apps_files/operation_artifact.txt
```

1101 3.5.7 Import definition

1102 An import definition is used within a TOSCA Service Template to locate and uniquely name another
1103 TOSCA Service Template file which has type and template definitions to be imported (included) and
1104 referenced within another Service Template.

1105 3.5.7.1 Keynames

1106 The following is the list of recognized keynames for a TOSCA import definition:

Keyname	Required	Type	Constraints	Description
file	yes	string	None	The required symbolic name for the imported file.
repository	no	string	None	The optional symbolic name of the repository definition where the imported file can be found as a string.
namespace_uri	no	string	None	The optional namespace URI to that will be applied to type definitions found within the imported file as a string.
namespace_prefix	no	string	None	The optional namespace prefix (alias) that will be used to indicate the namespace_uri when forming a qualified name (i.e., QName) when referencing type definitions from the imported file.

1107 3.5.7.2 Grammar

1108 Import definitions have one the following grammars:

1109 3.5.7.2.1 Single-line grammar:

```
imports:  
- <file_URI_1>  
- <file_URI_2>
```

1110 3.5.7.2.2 Multi-line grammar

```
imports:  
- file: <file_URI>  
  repository: <repository_name>  
  namespace_uri: <definition_namespace_uri>  
  namespace_prefix: <definition_namespace_prefix>
```

1111 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1112 • **file_uri**: contains the required name (i.e., URI) of the file to be imported as a [string](#).
- 1113 • **repository_name**: represents the optional symbolic name of the repository definition where the
1114 imported file can be found as a [string](#).
- 1115 • **namespace_uri**: represents the optional namespace URI to that will be applied to type
1116 definitions found within the imported file as a [string](#).
- 1117 • **namespace_prefix**: represents the optional namespace prefix (alias) that will be used to
1118 indicate the **namespace_uri** when forming a qualified name (i.e., QName) when referencing type
1119 definitions from the imported file as a [string](#).

1120 3.5.7.3 Example

1121 The following represents how import definitions would be used for the imports keyname within a TOSCA
1122 Service Template:

```
imports:
```

```

- some_definition_file: path1/path2/some_defs.yaml
- another_definition_file:
  file: path1/path2/file2.yaml
  repository: my_service_catalog
  namespace_uri: http://mycompany.com/tosca/1.0/platform
  namespace_prefix: mycompany

```

1123 3.5.8 Property definition

1124 A property definition defines a named, typed value and related data that can be associated with an entity
 1125 defined in this specification (e.g., Node Types, Relationship Types, Capability Types, etc.). Properties
 1126 are used by template authors to provide input values to TOSCA entities which indicate their “desired
 1127 state” when they are instantiated. The value of a property can be retrieved using the **get_property**
 1128 function within TOSCA Service Templates.

1129 3.5.8.1.1 Attribute and Property reflection

1130 The actual state of the entity, at any point in its lifecycle once instantiated, is reflected by [Attribute](#)
 1131 [definitions](#). TOSCA orchestrators automatically create an attribute for every declared property (with the
 1132 same symbolic name) to allow introspection of both the desired state (property) and actual state
 1133 (attribute).

1134 3.5.8.2 Keynames

1135 The following is the list of recognized keynames for a TOSCA property definition:

Keyname	Required	Type	Constraints	Description
type	yes	string	None	The required data type for the property.
description	no	description	None	The optional description for the property.
required	no	boolean	default: true	An optional key that declares a property as required (true) or not (false).
default	no	<any>	None	An optional key that may provide a value to be used as a default if not provided by another means.
status	no	string	default: supported	The optional status of the property relative to the specification or implementation. See table below for valid values.
constraints	no	list of constraint clauses	None	The optional list of sequenced constraint clauses for the property.
entry_schema	no	string	None	The optional key that is used to declare the name of the Datatype definition for entries of set types such as the TOSCA list or map .

1136 3.5.8.3 Status values

1137 The following property status values are supported:

Value	Description
supported	Indicates the property is supported. This is the default value for all property definitions.
unsupported	Indicates the property is not supported.
experimental	Indicates the property is experimental and has no official standing.

Value	Description
deprecated	Indicates the property has been deprecated by a new specification version.

1138 3.5.8.4 Grammar

1139 Named property definitions have the following grammar:

```

<property_name>:
  type: <property_type>
  description: <property_description>
  required: <property_required>
  default: <default_value>
  status: <status_value>
  constraints:
    - <property_constraints>
  entry_schema:
    description: <entry_description>
    type: <entry_type>
    constraints:
      - <entry_constraints>

```

1140 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1141 • **property_name**: represents the required symbolic name of the property as a [string](#).
- 1142 • **property_description**: represents the optional [description](#) of the property.
- 1143 • **property_type**: represents the required data type of the property.
- 1144 • **property_required**: represents an optional [boolean](#) value (true or false) indicating whether or
- 1145 not the property is required. If this keyname is not present on a property definition, then the
- 1146 property SHALL be considered **required** (i.e., true) by **default**.
- 1147 • **default_value**: contains a type-compatible value that may be used as a default if not provided
- 1148 by another means.
- 1149 • **status_value**: a [string](#) that contains a keyword that indicates the status of the property relative
- 1150 to the specification or implementation.
- 1151 • **property_constraints**: represents the optional [sequenced](#) list of one or more [constraint](#)
- 1152 [clauses](#) on the property definition.
- 1153 • **entry_description**: represents the optional [description](#) of the entry schema.
- 1154 • **entry_type**: represents the required type name for entries in a [list](#) or [map](#) property type.
- 1155 • **entry_constraints**: represents the optional [sequenced](#) list of one or more [constraint clauses](#)
- 1156 on entries in a [list](#) or [map](#) property type.

1157 3.5.8.5 Additional Requirements

- 1158 • Implementations of the TOSCA Simple Profile **SHALL** automatically reflect (i.e., make available)
- 1159 any property defined on an entity as an attribute of the entity with the same name as the property.
- 1160 • A property **SHALL** be considered [required by default](#) (i.e., as if the **required** keyname on the
- 1161 definition is set to **true**) unless the definition's **required** keyname is explicitly set to **false**.
- 1162 • The value provided on a property definition's **default** keyname **SHALL** be type compatible with
- 1163 the type declared on the definition's **type** keyname.
- 1164 • Constraints of a property definition **SHALL** be type-compatible with the type defined for that
- 1165 definition.

1166 3.5.8.6 Notes

- 1167 • This element directly maps to the **PropertiesDefinition** element defined as part of the
1168 schema for most type and entities defined in the [TOSCA v1.0 specification](#).
- 1169 • In the [TOSCA v1.0 specification](#) constraints are expressed in the XML Schema definitions of
1170 Node Type properties referenced in the **PropertiesDefinition** element of **NodeType**
1171 definitions.

1172 3.5.8.7 Example

1173 The following represents an example of a property definition with constraints:

```
properties:  
  num_cpus:  
    type: integer  
    description: Number of CPUs requested for a software node instance.  
    default: 1  
    required: true  
    constraints:  
      - valid_values: [ 1, 2, 4, 8 ]
```

1174 3.5.9 Property assignment

1175 This section defines the grammar for assigning values to named properties within TOSCA Node and
1176 Relationship templates that are defined in their corresponding named types.

1177 3.5.9.1 Keynames

1178 The TOSCA property assignment has no keynames.

1179 3.5.9.2 Grammar

1180 Property assignments have the following grammar:

1181 3.5.9.2.1 Short notation:

1182 The following single-line grammar may be used when a simple value assignment is needed:

```
<property_name>: <property_value> | { <property_value_expression> }
```

1183 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1184 • **property_name**: represents the name of a property that would be used to select a property
1185 definition with the same name within on a TOSCA entity (e.g., Node Template, Relationship
1186 Template, etc.) which is declared in its declared type (e.g., a Node Type, Node Template,
1187 Capability Type, etc.).
- 1188 • **property_value**, **property_value_expression**: represent the type-compatible value to
1189 assign to the named property. Property values may be provided as the result from the
1190 evaluation of an expression or a function.

1191 3.5.10 Attribute definition

1192 An attribute definition defines a named, typed value that can be associated with an entity defined in this
1193 specification (e.g., a Node, Relationship or Capability Type). Specifically, it is used to expose the “actual
1194 state” of some property of a TOSCA entity after it has been deployed and instantiated (as set by the
1195 TOSCA orchestrator). Attribute values can be retrieved via the **get_attribute** function from the
1196 instance model and used as values to other entities within TOSCA Service Templates.

1197 **3.5.10.1 Attribute and Property reflection**

1198 TOSCA orchestrators automatically create [Attribute definitions](#) for any [Property definitions](#) declared on
1199 the same TOSCA entity (e.g., nodes, node capabilities and relationships) in order to make accessible the
1200 actual (i.e., the current state) value from the running instance of the entity.

1201 **3.5.10.2 Keynames**

1202 The following is the list of recognized keynames for a TOSCA attribute definition:

Keyname	Required	Type	Constraints	Description
type	yes	string	None	The required data type for the attribute.
description	no	description	None	The optional description for the attribute.
default	no	<any>	None	An optional key that may provide a value to be used as a default if not provided by another means. This value SHALL be type compatible with the type declared by the property definition's type keyname.
status	no	string	default: supported	The optional status of the attribute relative to the specification or implementation. See supported status values defined under the Property definition section.
entry_schema	no	string	None	The optional key that is used to declare the name of the Datatype definition for entries of set types such as the TOSCA list or map .

1203 **3.5.10.3 Grammar**

1204 Attribute definitions have the following grammar:

```
attributes:  
  <attribute_name>:  
    type: <attribute_type>  
    description: <attribute_description>  
    default: <default_value>  
    status: <status_value>
```

1205 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1206 • **attribute_name**: represents the required symbolic name of the attribute as a [string](#).
- 1207 • **attribute_type**: represents the required data type of the attribute.
- 1208 • **attribute_description**: represents the optional [description](#) of the attribute.
- 1209 • **default_value**: contains a type-compatible value that may be used as a default if not provided
1210 by another means.
- 1211 • **status_value**: contains a value indicating the attribute's status relative to the specification
1212 version (e.g., supported, deprecated, etc.). Supported [status values](#) for this keyname are defined
1213 under [Property definition](#).

1214 **3.5.10.4 Additional Requirements**

- 1215 • In addition to any explicitly defined attributes on a TOSCA entity (e.g., Node Type,
1216 RelationshipType, etc.), implementations of the TOSCA Simple Profile **MUST** automatically
1217 reflect (i.e., make available) any property defined on an entity as an attribute of the entity with the
1218 same name as the property.

- 1219
- 1220
- 1221
- Values for the default keyname **MUST** be derived or calculated from other attribute or operation output values (that reflect the actual state of the instance of the corresponding resource) and not hard-coded or derived from a property settings or inputs (i.e., desired state).

1222

3.5.10.5 Notes

- 1223
- 1224
- 1225
- 1226
- 1227
- 1228
- 1229
- Attribute definitions are very similar to [Property definitions](#); however, properties of entities reflect an input that carries the template author's requested or desired value (i.e., desired state) which the orchestrator (attempts to) use when instantiating the entity whereas attributes reflect the actual value (i.e., actual state) that provides the actual instantiated value.
 - For example, a property can be used to request the IP address of a node using a property (setting); however, the actual IP address after the node is instantiated may be different and made available by an attribute.

1230

3.5.10.6 Example

1231 The following represents a required attribute definition:

```
actual_cpus:  
  type: integer  
  description: Actual number of CPUs allocated to the node instance.
```

1232

3.5.11 Attribute assignment

1233 This section defines the grammar for assigning values to named attributes within TOSCA Node and
1234 Relationship templates which are defined in their corresponding named types.

1235

3.5.11.1 Keynames

1236 The TOSCA attribute assignment has no keynames.

1237

3.5.11.2 Grammar

1238 Attribute assignments have the following grammar:

1239

3.5.11.2.1 Short notation:

1240 The following single-line grammar may be used when a simple value assignment is needed:

```
<attribute_name>: <attribute_value> | { <attribute_value_expression> }
```

1241

3.5.11.2.2 Extended notation:

1242 The following multi-line grammar may be used when a value assignment requires keys in addition to a
1243 simple value assignment:

```
<attribute_name>:  
  description: <attribute_description>  
  value: <attribute_value> | { <attribute_value_expression> }
```

1244 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1245
- 1246
- 1247
- 1248
- **attribute_name**: represents the name of an attribute that would be used to select an attribute definition with the same name within on a TOSCA entity (e.g., Node Template, Relationship Template, etc.) which is declared (or reflected from a Property definition) in its declared type (e.g., a Node Type, Node Template, Capability Type, etc.).

- 1249 • **attribute_value, attribute_value_expresssion:** represent the type-compatible value to
1250 assign to the named attribute. Attribute values may be provided as the result from the
1251 evaluation of an expression or a function.
- 1252 • **attribute_description:** represents the optional [description](#) of the attribute.

1253 3.5.11.3 Additional requirements

- 1254 • Attribute values **MAY** be provided by the underlying implementation at runtime when requested
1255 by the `get_attribute` function or it **MAY** be provided through the evaluation of expressions and/or
1256 functions that derive the values from other TOSCA attributes (also at runtime).

1257 3.5.12 Parameter definition

1258 A parameter definition is essentially a TOSCA property definition; however, it also allows a value to be
1259 assigned to it (as for a TOSCA property assignment). In addition, in the case of output parameters, it can
1260 optionally inherit the data type of the value assigned to it rather than have an explicit data type defined for
1261 it.

1262 3.5.12.1 Keynames

1263 The TOSCA parameter definition has all the keynames of a TOSCA Property definition, but in addition
1264 includes the following additional or changed keynames:

Keyname	Required	Type	Constraints	Description
type	no	string	None	The required data type for the parameter. Note: This keyname is required for a TOSCA Property definition, but is not for a TOSCA Parameter definition.
value	no	<any>	N/A	The type-compatible value to assign to the named parameter. Parameter values may be provided as the result from the evaluation of an expression or a function.

1265 3.5.12.2 Grammar

1266 Named parameter definitions have the following grammar:

```

<parameter_name>:
  type: <parameter_type>
  description: <parameter_description>
  value: <parameter_value> | { <parameter_value_expression> }
  required: <parameter_required>
  default: <parameter_default_value>
  status: <status_value>
  constraints:
    - <parameter_constraints>
  entry_schema:
    description: <entry_description>
    type: <entry_type>
    constraints:
      - <entry_constraints>

```

1267 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1268 • **parameter_name:** represents the required symbolic name of the parameter as a [string](#).
- 1269 • **parameter_description:** represents the optional [description](#) of the parameter.

- 1270 • **parameter_type**: represents the optional data type of the parameter. Note, this keyname is
1271 required for a TOSCA Property definition, but is not for a TOSCA Parameter definition.
- 1272 • **parameter_value**, **parameter_value_expresssion**: represent the type-compatible value to
1273 assign to the named parameter. Parameter values may be provided as the result from the
1274 evaluation of an expression or a function.
- 1275 • **parameter_required**: represents an optional **boolean** value (true or false) indicating whether or
1276 not the parameter is required. If this keyname is not present on a parameter definition, then the
1277 property SHALL be considered **required** (i.e., true) by **default**.
- 1278 • **default_value**: contains a type-compatible value that may be used as a default if not provided
1279 by another means.
- 1280 • **status_value**: a **string** that contains a keyword that indicates the status of the parameter
1281 relative to the specification or implementation.
- 1282 • **parameter_constraints**: represents the optional *sequenced* list of one or more **constraint**
1283 **clauses** on the parameter definition.
- 1284 • **entry_description**: represents the optional **description** of the entry schema.
- 1285 • **entry_type**: represents the required type name for entries in a **list** or **map** parameter type.
- 1286 • **entry_constraints**: represents the optional *sequenced* list of one or more **constraint clauses**
1287 on entries in a **list** or **map** parameter type.

1288 3.5.12.3 Additional Requirements

- 1289 • A parameter **SHALL** be considered *required by default* (i.e., as if the **required** keyname on the
1290 definition is set to **true**) unless the definition's **required** keyname is explicitly set to **false**.
- 1291 • The value provided on a parameter definition's **default** keyname **SHALL** be type compatible
1292 with the type declared on the definition's **type** keyname.
- 1293 • Constraints of a parameter definition **SHALL** be type-compatible with the type defined for that
1294 definition.

1295 3.5.12.4 Example

1296 The following represents an example of an input parameter definition with constraints:

```
inputs:
  cpus:
    type: integer
    description: Number of CPUs for the server.
    constraints:
      - valid_values: [ 1, 2, 4, 8 ]
```

1297 The following represents an example of an (untyped) output parameter definition:

```
outputs:
  server_ip:
    description: The private IP address of the provisioned server.
    value: { get_attribute: [ my_server, private_address ] }
```

1298

1299 3.5.13 Operation definition

1300 An operation definition defines a named function or procedure that can be bound to an implementation
1301 artifact (e.g., a script).

1302 3.5.13.1 Keynames

1303 The following is the list of recognized keynames for a TOSCA operation definition:

Keyname	Required	Type	Description
description	no	description	The optional description string for the associated named operation.
implementation	no	string	The optional implementation artifact name (e.g., a script file name within a TOSCA CSAR file).
inputs	no	list of property definitions	The optional list of input properties definitions (i.e., parameter definitions) for operation definitions that are within TOSCA Node or Relationship Type definitions. This includes when operation definitions are included as part of a Requirement definition in a Node Type.
	no	list of property assignments	The optional list of input property assignments (i.e., parameters assignments) for operation definitions that are within TOSCA Node or Relationship Template definitions. This includes when operation definitions are included as part of a Requirement assignment in a Node Template.

1304 The following is the list of recognized keynames to be used with the **implementation** keyname within a
1305 TOSCA operation definition:

Keyname	Required	Type	Description
primary	no	string	The optional implementation artifact name (i.e., the primary script file name within a TOSCA CSAR file).
dependencies	no	list of string	The optional ordered list of one or more dependent or secondary implementation artifact name which are referenced by the primary implementation artifact (e.g., a library the script installs or a secondary script).

1306 3.5.13.2 Grammar

1307 Operation definitions have the following grammars:

1308 3.5.13.2.1 Short notation

1309 The following single-line grammar may be used when only an operation's implementation artifact is
1310 needed:

```
<operation_name>: <implementation artifact name>
```

1311 3.5.13.2.2 Extended notation for use in Type definitions

1312 The following multi-line grammar may be used in Node or Relationship Type definitions when additional
1313 information about the operation is needed:

```
<operation_name>:
  description: <operation description>
  implementation: <implementation artifact name>
  inputs:
    <property definitions>
```

1314 3.5.13.2.3 Extended notation for use in Template definitions

1315 The following multi-line grammar may be used in Node or Relationship Template definitions when there
1316 are multiple artifacts that may be needed for the operation to be implemented:


```
<operation_name>:
  description: <operation_description>
  implementation:
    primary: <implementation_artifact_name>
    dependencies:
      - <list_of_dependent_artifact_names>
  inputs:
    <property_assignments>
```

1317 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1318 • **operation_name**: represents the required symbolic name of the operation as a [string](#).
- 1319 • **operation_description**: represents the optional [description](#) string for the corresponding
- 1320 **operation_name**.
- 1321 • **implementation_artifact_name**: represents the optional name ([string](#)) of an implementation
- 1322 artifact definition (defined elsewhere), or the direct name of an implementation artifact's relative
- 1323 filename (e.g., a service template-relative, path-inclusive filename or absolute file location using a
- 1324 URL).
- 1325 • **property_definitions**: represents the optional list of [property definitions](#) which the TOSCA
- 1326 orchestrator would make available (i.e., or pass) to the corresponding implementation artifact
- 1327 during its execution.
- 1328 • **property_assignments**: represents the optional list of [property assignments](#) for passing
- 1329 parameters to Node or Relationship Template operations providing values for properties defined
- 1330 in their respective type definitions.
- 1331 • **list_of_dependent_artifact_names**: represents the optional ordered list of one or more
- 1332 dependent or secondary implementation artifact names (as strings) which are referenced by the
- 1333 primary implementation artifact. TOSCA orchestrators will copy these files to the same location
- 1334 as the primary artifact on the target node so as to make them accessible to the primary
- 1335 implementation artifact when it is executed.

1336 3.5.13.3 Additional requirements

- 1337 • The default sub-classing behavior for implementations of operations SHALL be override. That is,
- 1338 implementation artifacts assigned in subclasses override any defined in its parent class.
- 1339 • Template authors MAY provide property assignments on operation inputs on templates that do
- 1340 not necessarily have a property definition defined in its corresponding type.
- 1341 • Implementation artifact file names (e.g., script filenames) may include file directory path names
- 1342 that are relative to the TOSCA service template file itself when packaged within a TOSCA Cloud
- 1343 Service ARchive (CSAR) file.

1344 3.5.13.4 Examples

1345 3.5.13.4.1 Single-line implementation example

```
interfaces:
  Standard:
    start: scripts/start_server.sh
```

1346 3.5.13.4.2 Multi-line implementation example

```
interfaces:
  Configure:
```



```

pre_configure_source:
  implementation:
    primary: scripts/pre_configure_source.sh
  dependencies:
    - scripts/setup.sh
    - binaries/library.rpm
    - scripts/register.py

```

1347 3.5.14 Interface definition

1348 An interface definition defines a named interface that can be associated with a Node or Relationship Type

1349 3.5.14.1 Keynames

1350 The following is the list of recognized keynames for a TOSCA interface definition:

Keyname	Required	Type	Description
inputs	no	list of property definitions	The optional list of input property definitions available to all defined operations for interface definitions that are within TOSCA Node or Relationship Type definitions. This includes when interface definitions are included as part of a Requirement definition in a Node Type.
	no	list of property assignments	The optional list of input property assignments (i.e., parameters assignments) for interface definitions that are within TOSCA Node or Relationship Template definitions. This includes when interface definitions are referenced as part of a Requirement assignment in a Node Template.

1351 3.5.14.2 Grammar

1352 Interface definitions have the following grammar:

1353 3.5.14.2.1 Extended notation for use in Type definitions

1354 The following multi-line grammar may be used in Node or Relationship Type definitions:

```

<interface definition name>:
  type: <interface type name>
  inputs:
    <property definitions>
    <operation definitions>

```

1355 3.5.14.2.2 Extended notation for use in Template definitions

1356 The following multi-line grammar may be used in Node or Relationship Template definitions:

```

<interface definition name>:
  inputs:
    <property assignments>
    <operation definitions>

```

1357 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1358 • **interface_definition_name:** represents the required symbolic name of the interface as a
- 1359 [string](#).
- 1360 • **interface_type_name:** represents the required name of the Interface Type for the interface
- 1361 **definition.**

- 1362 • **property_definitions**: represents the optional list of [property definitions](#) (i.e., parameters)
1363 which the TOSCA orchestrator would make available (i.e., or pass) to all defined operations.
1364 - *This means these properties and their values would be accessible to the implementation*
1365 *artifacts (e.g., scripts) associated to each operation during their execution.*
- 1366 • **property_assignments**: represents the optional list of [property assignments](#) for passing
1367 parameters to Node or Relationship Template operations providing values for properties defined
1368 in their respective type definitions.
- 1369 • **operation_definitions**: represents the required name of one or more [operation definitions](#).

1370 3.5.15 Event Filter definition

1371 An event filter definition defines criteria for selection of an attribute, for the purpose of monitoring it, within
1372 a TOSCA entity, or one its capabilities.

1373 3.5.15.1 Keynames

1374 The following is the list of recognized keynames for a TOSCA event filter definition:

Keyname	Required	Type	Description
node	yes	string	The required name of the node type or template that contains either the attribute to be monitored or contains the requirement that references the node that contains the attribute to be monitored.
requirement	no	string	The optional name of the requirement within the filter's node that can be used to locate a referenced node that contains an attribute to monitor.
capability	no	string	The optional name of a capability within the filter's node or within the node referenced by its requirement that contains the attribute to monitor.

1375 3.5.15.2 Grammar

1376 Event filter definitions have following grammar:

```
node: <node_type_name> | <node_template_name>
requirement: <requirement_name>
capability: <capability_name>
```

1377 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1378 • **node_type_name**: represents the required name of the node type that would be used to select
1379 (filter) the node that contains the attribute to monitor or contains the requirement that references
1380 another node that contains the attribute to monitor.
- 1381 • **node_template_name**: represents the required name of the node template that would be used to
1382 select (filter) the node that contains the attribute to monitor or contains the requirement that
1383 references another node that contains the attribute to monitor.
- 1384 • **requirement_name**: represents the optional name of the requirement that would be used to
1385 select (filter) a referenced node that contains the attribute to monitor.
- 1386 • **capability_name**: represents the optional name of a capability that would be used to select
1387 (filter) the attribute to monitor.

1388 3.5.16 Trigger definition

1389 A trigger definition defines the event, condition and action that is used to “trigger” a policy it is associated
1390 with.

1391 3.5.16.1 Keynames

1392 The following is the list of recognized keynames for a TOSCA trigger definition:

Keyname	Required	Type	Description
description	no	description	The optional description string for the named trigger.
event_type	yes	string	The required name of the event type that activates the trigger’s action.
schedule	no	TimeInterval	The optional time interval during which the trigger is valid (i.e., during which the declared actions will be processed).
target_filter	no	event filter	The optional filter used to locate the attribute to monitor for the trigger’s defined condition. This filter helps locate the TOSCA entity (i.e., node or relationship) or further a specific capability of that entity that contains the attribute to monitor.
condition	no	constraint clause	The optional condition which contains an attribute constraint that can be monitored. Note: this is optional since sometimes the event occurrence itself is enough to trigger the action.
constraint	no	constraint clause	The optional condition which contains an attribute constraint that can be monitored. Note: this is optional since sometimes the event occurrence itself is enough to trigger the action.
period	no	scalar-unit.time	The optional period to use to evaluate for the condition.
evaluations	no	integer	The optional number of evaluations that must be performed over the period to assert the condition exists.
method	no	string	The optional statistical method name to use to perform the evaluation of the condition.
action	yes	string or operation	The if of the workflow to be invoked when the event is triggered and the condition is met (i.e, evaluates to true). Or The required operation to invoke when the event is triggered and the condition is met (i.e., evaluates to true).

1393 3.5.16.2 Grammar

1394 Trigger definitions have the following grammars:

```
<trigger_name>:  
  description: <trigger_description>  
  # TBD: need to separate “simple” and “full” grammar for event type name  
  event: <event_type_name>  
    type: <event_type_name>  
  schedule: <time_interval_for_trigger>  
  target_filter:  
    <event_filter_definition>  
  condition: <attribute_constraint_clause>  
    constraint: <constraint_clause>  
    period: <scalar-unit.time> # e.g., 60 sec  
    evaluations: <integer> # e.g., 1  
    method: <string> # e.g., average  
  action:  
    <operation_definition>
```

1395 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1396 • **trigger_name**: represents the required symbolic name of the trigger as a [string](#).
- 1397 • **trigger_description**: represents the optional [description](#) string for the corresponding
- 1398 **trigger_name**.
- 1399 • **event_type_name**: represents the required name of the TOSCA Event Type that would be
- 1400 monitored on the identified resource (node).
- 1401 • **time_interval_for_trigger**: **represents the optional time interval that the trigger is valid**
- 1402 **for**.
- 1403 • **event_filter_definition**: represents the optional filter to use to locate the resource (node)
- 1404 or capability attribute to monitor.
- 1405 • **attribute_constraint_clause**: represents the optional attribute constraint that would be
- 1406 used to test for a specific condition on the monitored resource.
- 1407 • **operation_definition**: represents the required action to take if the event and (optionally)
- 1408 condition are met.

1409 3.5.17 Workflow activity definition

1410 A workflow activity defines an operation to be performed in a TOSCA workflow. Activities allows to:

1411

- 1412 • Delegate the workflow for a node expected to be provided by the orchestrator
- 1413 • Set the state of a node
- 1414 • Call an operation defined on a TOSCA interface of a node, relationship or group
- 1415 • Inline another workflow defined in the topology (to allow reusability)

1416 3.5.17.1 Keynames

1417 The following is the list of recognized keynames for a TOSCA workflow activity definition. Note that while
1418 each of the key is not required, one and only one of them is required (mutually exclusive).

Keyname	Required	Type	Description
delegate	no	string	The name of the delegate workflow. This activity requires the target to be provided by the orchestrator (no-op node or relationship)
set_state	no	string	Value of the node state.
call_operation	no	string	A string that defines the name of the interface and operation to be called on the node using the <interface_name>.<operation_name> notation.
inline	no	string	The name of a workflow to be inlined.

1419 3.5.17.2 Grammar

1420 Workflow activity definitions have one of the following grammars:

1421 3.5.17.2.1 Delegate activity

```
- delegate: <delegate_workflow_name>
```

1422 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1423 • **delegate_workflow_name:** represents the name of the workflow of the node
1424 provided by the TOSCA orchestrator.

1425 **3.5.17.2.2 Set state activity**

```
- set_state: <new_node_state>
```

1426 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1427 • **new_node_state:** represents the state that will be affected to the node once
1428 the activity is performed.

1429 **3.5.17.2.3 Call operation activity:**

```
- call_operation: <interface_name>.<operation_name>
```

1430 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1431 • **interface_name:** represents the name of the interface in which the operation to
1432 be called is defined.
1433 • **operation_name:** represents the name of the operation of the interface that
1434 will be called during the workflow execution.

1435 **3.5.17.2.4 Inline activity**

```
- inline: <workflow_name>
```

1436 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1437 • **workflow_name:** represents the name of the workflow to inline.

1438 **3.5.17.3 Additional Requirements**

- 1439 • Keynames are mutually exclusive, i.e. an activity **MUST** define only one of `delegate`, `set_state`,
1440 `call_operation` or `inline` keyname.

1441 **3.5.17.4 Example**

1442 following represents a list of workflow activity definitions:

```
- delegate: deploy  
- set_state: started  
- call_operation: tosca.interfaces.node.lifecycle.Standard.start  
- inline: my_workflow
```

1443

1444 **3.5.18 Assertion definition**

1445 A workflow assertion is used to specify a single condition on a workflow filter definition. The assertion
1446 allows to assert the value of an attribute based on TOSCA constraints.

1447 **3.5.18.1 Keynames**

1448 The TOSCA workflow assertion definition has no keynames.

1449 **3.5.18.2 Grammar**

1450 Workflow assertion definitions have the following grammar:

```
<attribute_name>: <list_of_constraint_clauses>
```

1451 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1452 • **attribute_name**: represents the name of an attribute defined on the assertion context entity
1453 (node instance, relationship instance, group instance) and from which value will be evaluated
1454 against the defined constraint clauses.
- 1455 • **list_of_constraint_clauses**: represents the list of constraint clauses that will be used to validate
1456 the attribute assertion.

1457 3.5.18.3 Example

1458 Following represents a workflow assertion with a single equals constraint:

```
my_attribute: [{equal : my_value}]
```

1459 Following represents a workflow assertion with multiple constraints:

```
my_attribute:  
- min_length: 8  
- max_length : 10
```

1460 3.5.19 Condition clause definition

1461 A workflow condition clause definition is used to specify a condition that can be used within a workflow
1462 precondition or workflow filter.

1463 3.5.19.1 Keynames

1464 The following is the list of recognized keynames for a TOSCA workflow condition definition:

Keyname	Required	Type	Description
and	no	list of condition clause definition	An and clause allows to define sub-filter clause definitions that must all be evaluated truly so the and clause is considered as true.
or	no	list of condition clause definition	An or clause allows to define sub-filter clause definitions where one of them must all be evaluated truly so the or clause is considered as true. Note in opposite to assert
assert	no	list of assertion definition	A list of filter assertions to be evaluated on entity attributes. Assert acts as a and clause, i.e. every defined filter assertion must be true so the assertion is considered as true.

1465 3.5.19.2 Grammar

1466 Workflow assertion definitions have the following grammars:

1467 3.5.19.2.1 And clause

```
and: <list_of_condition_clause_definition>
```

1468 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1469 • **list_of_condition_clause_definition**: represents the list of condition clauses. All
1470 condition clauses MUST be asserted to true so that the and clause is asserted to true.

1471 **3.5.19.2.2 Or clause**

```
or: <list_of_condition_clause_definition>
```

1472 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1473
- `list_of_condition_clause_definition`: represents the list of condition clauses. One of the condition clause have to be asserted to true so that the or clause is asserted to true.
- 1474

1475 **3.5.19.2.3 Assert clause**

```
assert: <list_of_assertion_definition>
```

1476 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1477
- `list_of_assertion_definition`: represents the list of assertions. All assertions MUST be asserted to true so that the assert clause is asserted to true.
- 1478

1479 **3.5.19.3 Additional Requirement**

- 1480
- Keynames are mutually exclusive, i.e. a filter definition can define only one of *and*, *or*, or *assert* keyname.
- 1481

1482 **3.5.19.4 Notes**

- 1483
- The TOSCA processor SHOULD perform assertion in the order of the list for every defined condition clause or assertion definition.
- 1484

1485 **3.5.19.5 Example**

1486 Following represents a workflow condition clause with a single equals constraint:

```
condition:
- assert:
  - my_attribute: [{equal: my_value}]
```

1487 Following represents a workflow condition clause with a single equals constraints on two different attributes:

```
condition:
- assert:
  - my_attribute: [{equal: my_value}]
  - my_other_attribute: [{equal: my_other_value}]
```

1489 Following represents a workflow condition clause with a or constraint on two different assertions:

```
condition:
- or:
  - assert:
    - my_attribute: [{equal: my_value}]
  - assert:
    - my_other_attribute: [{equal: my_other_value}]
```

1490 Following represents multiple levels of condition clauses to build the following logic: one_attribute equal one_value AND (my_attribute equal my_value OR my_other_attribute equal my_other_value):

```
condition:
- assert:
```

```

- one_attribute: [{equal: one_value }]
- or:
- assert:
  - my_attribute: [{equal: my_value}]
- assert:
  - my_other_attribute: [{equal: my_other_value}]

```

1492 3.5.20 Workflow precondition definition

1493 A workflow condition can be used as a filter or precondition to check if a workflow can be processed or
 1494 not based on the state of the instances of a TOSCA topology deployment. When not met, the workflow
 1495 will not be triggered.

1496 3.5.20.1 Keynames

1497 The following is the list of recognized keynames for a TOSCA workflow condition definition:

Keyname	Required	Type	Description
target	yes	string	The target of the precondition (this can be a node template name, a group name)
target_relationship	no	string	The optional name of a requirement of the target in case the precondition has to be processed on a relationship rather than a node or group. Note that this is applicable only if the target is a node.
condition	no	list of condition clause definitions	A list of workflow condition clause definitions. Assertion between elements of the condition are evaluated as an AND condition.

1498 3.5.20.2 Grammar

1499 Workflow precondition definitions have the following grammars:

```

- target: <target_name>
  target_relationship: <target_requirement_name>
  condition:
    <list_of_condition_clause_definition>

```

1500 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1501 • **target_name**: represents the name of a node template or group in the topology.
- 1502 • **target_requirement_name**: represents the name of a requirement of the node template (in case
 1503 target_name refers to a node template.
- 1504 • **list_of_condition_clause_definition**: represents the list of condition clauses
 1505 to be evaluated. The value of the resulting condition is evaluated as an AND
 1506 clause between the different elements.

1507 3.5.21 Workflow step definition

1508 A workflow step allows to define one or multiple sequenced activities in a workflow and how they are
 1509 connected to other steps in the workflow. They are the building blocks of a declarative workflow.

1510 3.5.21.1 Keynames

1511 The following is the list of recognized keynames for a TOSCA workflow step definition:

Keyname	Required	Type	Description
target	yes	string	The target of the step (this can be a node template name, a group name)
target_relationship	no	string	The optional name of a requirement of the target in case the step refers to a relationship rather than a node or group. Note that this is applicable only if the target is a node.
operation_host	no	string	The node on which operations should be executed (for TOSCA call_operation activities). This element is required only for relationships and groups target. If target is a relationships operation_host is required and valid_values are SOURCE or TARGET – referring to the relationship source or target node. If target is a group operation_host is optional. If not specified the operation will be triggered on every node of the group. If specified the valid_value is a node_type or the name of a node template.
filter	no	list of constraint clauses	Filter is a map of attribute name, list of constraint clause that allows to provide a filtering logic.
activities	yes	list of activity_definition	The list of sequential activities to be performed in this step.
on_success	no	list of string	The optional list of step names to be performed after this one has been completed with success (all activities has been correctly processed).
on_failure	no	list of string	The optional list of step names to be called after this one in case one of the step activity failed.

1512 **3.5.21.2 Grammar**

1513 Workflow step definitions have the following grammars:

```

steps:
  <step_name>
    target: <target_name>
    target_relationship: <target_requirement_name>
    operation_host: <operation_host_name>
    filter:
      - <list_of_condition_clause_definition>
    activities:
      - <list_of_activity_definition>
    on_success:
      - <target_step_name>
    on_failure:
      - <target_step_name>

```

1514 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1515 • **target_name**: represents the name of a node template or group in the topology.
- 1516 • **target_requirement_name**: represents the name of a requirement of the node template (in case
- 1517 target_name refers to a node template.
- 1518 • **operation_host**: the node on which the operation should be executed
- 1519 • **<list_of_condition_clause_definition>**: represents a list of condition clause definition.

- 1520 • `list_of_activity_definition`: represents a list of activity definition
- 1521 • `target_step_name`: represents the name of another step of the workflow.

1522 3.6 Type-specific definitions

1523 3.6.1 Entity Type Schema

1524 An Entity Type is the common, base, polymorphic schema type which is extended by TOSCA base entity
 1525 type schemas (e.g., Node Type, Relationship Type, Artifact Type, etc.) and serves to define once all the
 1526 commonly shared keynames and their types. This is a “meta” type which is abstract and not directly
 1527 instantiatable.

1528 3.6.1.1 Keynames

1529 The following is the list of recognized keynames for a TOSCA Entity Type definition:

Keyname	Required	Type	Constraints	Description
derived_from	no	string	‘None’ is the only allowed value	An optional parent Entity Type name the Entity Type derives from.
version	no	version	N/A	An optional version for the Entity Type definition.
metadata	no	map of string	N/A	Defines a section used to declare additional metadata information.
description	no	description	N/A	An optional description for the Entity Type.

1530 3.6.1.2 Grammar

1531 Entity Types have following grammar:

```

<entity_keyname>:
  # The only allowed value is ‘None’
  derived_from: None
  version: <version_number>
  metadata:
    <metadata_map>
  description: <description>
  
```

1532 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1533 • `version_number`: represents the optional TOSCA `version` number for the entity.
- 1534 • `entity_description`: represents the optional `description` string for the entity.
- 1535 • `metadata_map`: represents the optional map of string.

1536 3.6.1.3 Additional Requirements

- 1537 • The TOSCA Entity Type SHALL be the common base type used to derive all other top-level base
 1538 TOSCA Types.
- 1539 • The TOSCA Entity Type SHALL NOT be used to derive or create new base types apart from
 1540 those defined in this specification or a profile of this specification.

1541 3.6.2 Capability definition

1542 A capability definition defines a named, typed set of data that can be associated with Node Type or Node
 1543 Template to describe a transparent capability or feature of the software component the node describes.

1544 **3.6.2.1 Keynames**

1545 The following is the list of recognized keynames for a TOSCA capability definition:

Keyname	Required	Type	Constraints	Description
type	yes	string	N/A	The required name of the Capability Type the capability definition is based upon.
description	no	description	N/A	The optional description of the Capability definition.
properties	no	list of property definitions	N/A	An optional list of property definitions for the Capability definition.
attributes	no	list of attribute definitions	N/A	An optional list of attribute definitions for the Capability definition.
valid_source_types	no	string[]	N/A	An optional list of one or more valid names of Node Types that are supported as valid sources of any relationship established to the declared Capability Type.
occurrences	no	range of integer	implied default of [1,UNBOUNDED]	The optional minimum and maximum occurrences for the capability. By default, an exported Capability should allow at least one relationship to be formed with it with a maximum of UNBOUNDED relationships. Note: the keyword UNBOUNDED is also supported to represent any positive integer.

1546 **3.6.2.2 Grammar**

1547 Capability definitions have one of the following grammars:

1548 **3.6.2.2.1 Short notation**

1549 The following grammar may be used when only a list of capability definition names needs to be declared:

```
<capability_definition_name>: <capability_type>
```

1550 **3.6.2.2.2 Extended notation**

1551 The following multi-line grammar may be used when additional information on the capability definition is
1552 needed:

```
<capability_definition_name>:
  type: <capability_type>
  description: <capability_description>
  properties:
    <property_definitions>
  attributes:
    <attribute_definitions>
  valid_source_types: [ <node_type_names> ]
```

1553 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1554 • **capability_definition_name**: represents the symbolic name of the capability as a **string**.
- 1555 • **capability_type**: represents the required name of a **capability type** the capability definition is
1556 based upon.
- 1557 • **capability_description**: represents the optional **description** of the capability definition.

- 1558
- **property_definitions**: represents the optional list of [property definitions](#) for the capability definition.
- 1559
- **attribute_definitions**: represents the optional list of [attribute definitions](#) for the capability definition.
- 1560
- **node_type_names**: represents the optional list of one or more names of [Node Types](#) that the Capability definition supports as valid sources for a successful relationship to be established to itself.
- 1561
- 1562
- 1563
- 1564

1565 3.6.2.3 Examples

1566 The following examples show capability definitions in both simple and full forms:

1567 3.6.2.3.1 Simple notation example

```
# Simple notation, no properties defined or augmented
some_capability: mytypes.mycapabilities.MyCapabilityTypeName
```

1568 3.6.2.3.2 Full notation example

```
# Full notation, augmenting properties of the referenced capability type
some_capability:
  type: mytypes.mycapabilities.MyCapabilityTypeName
  properties:
    limit:
      type: integer
      default: 100
```

1569 3.6.2.4 Additional requirements

- 1570
- Any Node Type (names) provides as values for the **valid_source_types** keyname SHALL be type-compatible (i.e., derived from the same parent Node Type) with any Node Types defined using the same keyname in the parent Capability Type.
- 1571
- Capability symbolic names SHALL be unique; it is an error if a capability name is found to occur more than once.
- 1572
- 1573
- 1574

1575 3.6.2.5 Notes

- 1576
- The Capability Type, in this example **MyCapabilityTypeName**, would be defined elsewhere and have an integer property named **limit**.
- 1577
- This definition directly maps to the **CapabilitiesDefinition** of the Node Type entity as defined in the [TOSCA v1.0 specification](#).
- 1578
- 1579

1580 3.6.3 Requirement definition

1581 The Requirement definition describes a named requirement (dependencies) of a TOSCA Node Type or Node template which needs to be fulfilled by a matching Capability definition declared by another TOSCA modelable entity. The requirement definition may itself include the specific name of the fulfilling entity (explicitly) or provide an abstract type, along with additional filtering characteristics, that a TOSCA orchestrator can use to fulfill the capability at runtime (implicitly).

1582

1583

1584

1585

1586 3.6.3.1 Keynames

1587 The following is the list of recognized keynames for a TOSCA requirement definition:

Keyname	Required	Type	Constraints	Description
capability	yes	string	N/A	The required reserved keyname used that can be used to provide the name of a valid Capability Type that can fulfill the requirement.
node	no	string	N/A	The optional reserved keyname used to provide the name of a valid Node Type that contains the capability definition that can be used to fulfill the requirement.
relationship	no	string	N/A	The optional reserved keyname used to provide the name of a valid Relationship Type to construct when fulfilling the requirement.
occurrences	no	range of integer	implied default of [1,1]	The optional minimum and maximum occurrences for the requirement. Note: the keyword UNBOUNDED is also supported to represent any positive integer.

1588 3.6.3.1.1 Additional Keynames for multi-line relationship grammar

1589 The Requirement definition contains the Relationship Type information needed by TOSCA Orchestrators
1590 to construct relationships to other TOSCA nodes with matching capabilities; however, it is sometimes
1591 recognized that additional properties may need to be passed to the relationship (perhaps for
1592 configuration). In these cases, additional grammar is provided so that the Node Type may declare
1593 additional Property definitions to be used as inputs to the Relationship Type's declared interfaces (or
1594 specific operations of those interfaces).

Keyname	Required	Type	Constraints	Description
type	yes	string	N/A	The optional reserved keyname used to provide the name of the Relationship Type for the requirement definition's relationship keyname.
interfaces	no	list of interface definitions	N/A	The optional reserved keyname used to reference declared (named) interface definitions of the corresponding Relationship Type in order to declare additional Property definitions for these interfaces or operations of these interfaces.

1595 3.6.3.2 Grammar

1596 Requirement definitions have one of the following grammars:

1597 3.6.3.2.1 Simple grammar (Capability Type only)

```
<requirement_name>: <capability_type_name>
```

1598 3.6.3.2.2 Extended grammar (with Node and Relationship Types)

```
<requirement_name>:
  capability: <capability_type_name>
  node: <node_type_name>
  relationship: <relationship_type_name>
  occurrences: [ <min_occurrences>, <max_occurrences> ]
```

1599 3.6.3.2.3 Extended grammar for declaring Property Definitions on the 1600 relationship's Interfaces

1601 The following additional multi-line grammar is provided for the relationship keyname in order to declare
1602 new Property definitions for inputs of known Interface definitions of the declared Relationship Type.

```
<requirement_name>:  
  # Other keynames omitted for brevity  
  relationship:  
    type: <relationship_type_name>  
    interfaces:  
      <interface_definitions>
```

1603 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1604 • **requirement_name**: represents the required symbolic name of the requirement definition as a
1605 [string](#).
- 1606 • **capability_type_name**: represents the required name of a Capability type that can be used to
1607 fulfill the requirement.
- 1608 • **node_type_name**: represents the optional name of a TOSCA Node Type that contains the
1609 Capability Type definition the requirement can be fulfilled by.
- 1610 • **relationship_type_name**: represents the optional name of a [Relationship Type](#) to be used to
1611 construct a relationship between this requirement definition (i.e., in the source node) to a
1612 matching capability definition (in a target node).
- 1613 • **min_occurrences**, **max_occurrences**: represents the optional minimum and maximum
1614 occurrences of the requirement (i.e., its cardinality).
- 1615 • **interface_definitions**: represents one or more already declared interface definitions in the
1616 Relationship Type (as declared on the **type** keyname) allowing for the declaration of new
1617 Property definition for these interfaces or for specific Operation definitions of these interfaces.

1618 3.6.3.3 Additional Requirements

- 1619 • Requirement symbolic names SHALL be unique; it is an error if a requirement name is found to
1620 occur more than once.
- 1621 • If the **occurrences** keyname is not present, then the occurrence of the requirement **SHALL** be
1622 one and only one; that is a default declaration as follows would be assumed:
1623 o occurrences: [1,1]

1624 3.6.3.4 Notes

- 1625 • This element directly maps to the **RequirementsDefinition** of the Node Type entity as defined
1626 in the [TOSCA v1.0 specification](#).
- 1627 • The requirement symbolic name is used for identification of the requirement definition only and
1628 not relied upon for establishing any relationships in the topology.

1629 3.6.3.5 Requirement Type definition is a tuple

1630 A requirement definition allows type designers to govern which types are allowed (valid) for fulfillment
1631 using three levels of specificity with only the Capability Type being required.

- 1632 1. Node Type (optional)
- 1633 2. Relationship Type (optional)
- 1634 3. Capability Type (required)

1635 The first level allows selection, as shown in both the simple or complex grammar, simply providing the
 1636 node's type using the **node** keyname. The second level allows specification of the relationship type to use
 1637 when connecting the requirement to the capability using the **relationship** keyname. Finally, the
 1638 specific named capability type on the target node is provided using the **capability** keyname.

1639 3.6.3.5.1 Property filter

1640 In addition to the node, relationship and capability types, a filter, with the keyname **node_filter**, may be
 1641 provided to constrain the allowed set of potential target nodes based upon their properties and their
 1642 capabilities' properties. This allows TOSCA orchestrators to help find the "best fit" when selecting among
 1643 multiple potential target nodes for the expressed requirements.

1644 3.6.4 Artifact Type

1645 An Artifact Type is a reusable entity that defines the type of one or more files that are used to define
 1646 implementation or deployment artifacts that are referenced by nodes or relationships on their operations.

1647 3.6.4.1 Keynames

1648 The Artifact Type is a TOSCA Entity and has the common keynames listed in section 3.6.1 TOSCA Entity
 1649 Schema.

1650 In addition, the Artifact Type has the following recognized keynames:

Keyname	Required	Type	Description
mime_type	no	string	The required mime type property for the Artifact Type.
file_ext	no	string[]	The required file extension property for the Artifact Type.
properties	no	list of property definitions	An optional list of property definitions for the Artifact Type.

1651 3.6.4.2 Grammar

1652 Artifact Types have following grammar:

```

<artifact_type_name>:
  derived_from: <parent_artifact_type_name>
  version: <version_number>
  metadata:
    <map of string>
  description: <artifact_description>
  mime_type: <mime_type_string>
  file_ext: [ <file_extensions> ]
  properties:
    <property_definitions>
  
```

1653 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1654 • **artifact_type_name**: represents the name of the Artifact Type being declared as a **string**.
- 1655 • **parent_artifact_type_name**: represents the **name** of the **Artifact Type** this Artifact Type
 1656 definition derives from (i.e., its "parent" type).
- 1657 • **version_number**: represents the optional TOSCA **version** number for the Artifact Type.
- 1658 • **artifact_description**: represents the optional **description** string for the Artifact Type.
- 1659 • **mime_type_string**: represents the optional Multipurpose Internet Mail Extensions (MIME)
 1660 standard string value that describes the file contents for this type of Artifact Type as a **string**.

- 1661 • **file_extensions**: represents the optional list of one or more recognized file extensions for this type of artifact type as [strings](#).
- 1662
- 1663 • **property_definitions**: represents the optional list of [property definitions](#) for the artifact type.

1664 3.6.4.3 Examples

```
my_artifact_type:
  description: Java Archive artifact type
  derived_from: tosca.artifact.Root
  mime_type: application/java-archive
  file_ext: [ jar ]
```

1665 3.6.4.4 Notes

- 1666 • The 'mime_type' keyname is meant to have values that are Apache mime types such as those defined here: <http://svn.apache.org/repos/asf/httpd/httpd/trunk/docs/conf/mime.types>
- 1667

1668 3.6.5 Interface Type

1669 An Interface Type is a reusable entity that describes a set of operations that can be used to interact with
1670 or manage a node or relationship in a TOSCA topology.

1671 3.6.5.1 Keynames

1672 The Interface Type is a TOSCA Entity and has the common keynames listed in section 3.6.1 TOSCA
1673 Entity Schema.

1674 In addition, the Interface Type has the following recognized keynames:

Keyname	Required	Type	Description
inputs	no	list of property definitions	The optional list of input parameter definitions.

1675 3.6.5.2 Grammar

1676 Interface Types have following grammar:

```
<interface type name>:
  derived_from: <parent interface type name>
  version: <version_number>
  metadata:
    <map of string>
  description: <interface description>
  inputs:
    <property definitions>
    <operation definitions>
```

1677 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1678 • **interface_type_name**: represents the required name of the interface as a [string](#).
- 1679 • **parent_interface_type_name**: represents the name of the [Interface Type](#) this Interface Type definition derives from (i.e., its "parent" type).
- 1680
- 1681 • **version_number**: represents the optional TOSCA [version](#) number for the Interface Type.
- 1682 • **interface_description**: represents the optional [description](#) string for the Interface Type.

- 1683 • **property_definitions**: represents the optional list of [property definitions](#) (i.e., parameters)
1684 which the TOSCA orchestrator would make available (i.e., or pass) to all implementation artifacts
1685 for operations declared on the interface during their execution.
- 1686 • **operation_definitions**: represents the required list of one or more [operation definitions](#).

1687 3.6.5.3 Example

1688 The following example shows a custom interface used to define multiple configure operations.

```
mycompany.mytypes.myinterfaces.MyConfigure:
  derived_from: tosca.interfaces.relationship.Root
  description: My custom configure Interface Type
  inputs:
    mode:
      type: string
  pre_configure_service:
    description: pre-configure operation for my service
  post_configure_service:
    description: post-configure operation for my service
```

1689 3.6.5.4 Additional Requirements

- 1690 • Interface Types **MUST NOT** include any implementations for defined operations; that is, the
1691 implementation keyname is invalid.
- 1692 • The **inputs** keyname is reserved and **SHALL NOT** be used for an operation name.

1693 3.6.6 Data Type

1694 A Data Type definition defines the schema for new named datatypes in TOSCA.

1695 3.6.6.1 Keynames

1696 The Data Type is a TOSCA Entity and has the common keynames listed in section 3.6.1 TOSCA Entity
1697 Schema.

1698 In addition, the Data Type has the following recognized keynames:

Keyname	Required	Type	Description
constraints	no	list of constraint clauses	The optional list of sequenced constraint clauses for the Data Type.
properties	no	list of property definitions	The optional list property definitions that comprise the schema for a complex Data Type in TOSCA.

1699 3.6.6.2 Grammar

1700 Data Types have the following grammar:

```
<data_type_name>:
  derived_from: <existing_type_name>
  version: <version_number>
  metadata:
    <map of string>
  description: <datatype_description>
  constraints:
```

```
- <type_constraints>
properties:
  <property_definitions>
```

1701 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1702 • **data_type_name**: represents the required symbolic name of the Data Type as a [string](#).
- 1703 • **version_number**: represents the optional TOSCA [version](#) number for the Data Type.
- 1704 • **datatype_description**: represents the optional [description](#) for the Data Type.
- 1705 • **existing_type_name**: represents the optional name of a valid TOSCA type this new Data
- 1706 Type would derive from.
- 1707 • **type_constraints**: represents the optional [sequenced](#) list of one or more type-compatible
- 1708 [constraint clauses](#) that restrict the Data Type.
- 1709 • **property_definitions**: represents the optional list of one or more [property definitions](#) that
- 1710 provide the schema for the Data Type.

1711 3.6.6.3 Additional Requirements

- 1712 • A valid datatype definition **MUST** have either a valid **derived_from** declaration or at least one
- 1713 valid property definition.
- 1714 • Any **constraint** clauses **SHALL** be type-compatible with the type declared by the
- 1715 **derived_from** keyname.
- 1716 • If a **properties** keyname is provided, it **SHALL** contain one or more valid property definitions.

1717 3.6.6.4 Examples

1718 The following example represents a Data Type definition based upon an existing string type:

1719 3.6.6.4.1 Defining a complex datatype

```
# define a new complex datatype
mytypes.phonenumber:
  description: my phone number datatype
  properties:
    countrycode:
      type: integer
    areacode:
      type: integer
    number:
      type: integer
```

1720 3.6.6.4.2 Defining a datatype derived from an existing datatype

```
# define a new datatype that derives from existing type and extends it
mytypes.phonenumber.extended:
  derived_from: mytypes.phonenumber
  description: custom phone number type that extends the basic phonenumber type
  properties:
    phone_description:
      type: string
    constraints:
      - max_length: 128
```

1721 3.6.7 Capability Type

1722 A Capability Type is a reusable entity that describes a kind of capability that a Node Type can declare to
1723 expose. Requirements (implicit or explicit) that are declared as part of one node can be matched to (i.e.,
1724 fulfilled by) the Capabilities declared by another node.

1725 3.6.7.1 Keynames

1726 The Capability Type is a TOSCA Entity and has the common keynames listed in section 3.6.1 TOSCA
1727 Entity Schema.

1728 In addition, the Capability Type has the following recognized keynames:

Keyname	Required	Type	Description
properties	no	list of property definitions	An optional list of property definitions for the Capability Type.
attributes	no	list of attribute definitions	An optional list of attribute definitions for the Capability Type.
valid_source_types	no	string []	An optional list of one or more valid names of Node Types that are supported as valid sources of any relationship established to the declared Capability Type.

1729 3.6.7.2 Grammar

1730 Capability Types have following grammar:

```
<capability type name>:  
  derived_from: <parent capability type name>  
  version: <version number>  
  description: <capability description>  
  properties:  
    <property definitions>  
  attributes:  
    <attribute definitions>  
  valid_source_types: [ <node type names> ]
```

1731 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1732 • **capability_type_name**: represents the required name of the Capability Type being declared as
1733 a [string](#).
- 1734 • **parent_capability_type_name**: represents the name of the [Capability Type](#) this Capability
1735 Type definition derives from (i.e., its “parent” type).
- 1736 • **version_number**: represents the optional TOSCA [version](#) number for the Capability Type.
- 1737 • **capability_description**: represents the optional [description](#) string for the corresponding
1738 **capability_type_name**.
- 1739 • **property_definitions**: represents an optional list of [property definitions](#) that the Capability
1740 type exports.
- 1741 • **attribute_definitions**: represents the optional list of [attribute definitions](#) for the Capability
1742 Type.
- 1743 • **node_type_names**: represents the optional list of one or more names of [Node Types](#) that the
1744 Capability Type supports as valid sources for a successful relationship to be established to itself.

1745 3.6.7.3 Example

```
mycompany.mytypes.myapplication.MyFeature:
  derived_from: tosca.capabilities.Root
  description: a custom feature of my company's application
  properties:
    my_feature_setting:
      type: string
    my_feature_value:
      type: integer
```

1746 3.6.8 Requirement Type

1747 A Requirement Type is a reusable entity that describes a kind of requirement that a Node Type can
1748 declare to expose. The TOSCA Simple Profile seeks to simplify the need for declaring specific
1749 Requirement Types from nodes and instead rely upon nodes declaring their features sets using TOSCA
1750 Capability Types along with a named Feature notation.

1751 Currently, there are no use cases in this TOSCA Simple Profile in YAML specification that utilize an
1752 independently defined Requirement Type. This is a desired effect as part of the simplification of the
1753 TOSCA v1.0 specification.

1754 3.6.9 Node Type

1755 A Node Type is a reusable entity that defines the type of one or more Node Templates. As such, a Node
1756 Type defines the structure of observable properties via a *Properties Definition, the Requirements and*
1757 *Capabilities of the node as well as its supported interfaces.*

1758 3.6.9.1 Keynames

1759 The Node Type is a TOSCA Entity and has the common keynames listed in section 3.6.1 TOSCA Entity
1760 Schema.

1761 In addition, the Node Type has the following recognized keynames:

Keyname	Required	Type	Description
attributes	no	list of attribute definitions	An optional list of attribute definitions for the Node Type.
properties	no	list of property definitions	An optional list of property definitions for the Node Type.
requirements	no	list of requirement definitions	An optional <i>sequenced</i> list of requirement definitions for the Node Type.
capabilities	no	list of capability definitions	An optional list of capability definitions for the Node Type.
interfaces	no	list of interface definitions	An optional list of interface definitions supported by the Node Type.
artifacts	no	list of artifact definitions	An optional list of named artifact definitions for the Node Type.

1762 3.6.9.2 Grammar

1763 Node Types have following grammar:

```
<node type name>:
```

```

derived_from: <parent_node_type_name>
version: <version_number>
metadata:
  <map of string>
description: <node_type_description>
attributes:
  <attribute_definitions>
properties:
  <property_definitions>
requirements:
  - <requirement_definitions>
capabilities:
  <capability_definitions>
interfaces:
  <interface_definitions>
artifacts:
  <artifact_definitions>

```

1764 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1765 • **node_type_name**: represents the required symbolic name of the Node Type being declared.
- 1766 • **parent_node_type_name**: represents the name ([string](#)) of the [Node Type](#) this Node Type
- 1767 definition derives from (i.e., its “parent” type).
- 1768 • **version_number**: represents the optional TOSCA [version](#) number for the Node Type.
- 1769 • **node_type_description**: represents the optional [description](#) string for the corresponding
- 1770 **node_type_name**.
- 1771 • **property_definitions**: represents the optional list of [property definitions](#) for the Node Type.
- 1772 • **attribute_definitions**: represents the optional list of [attribute definitions](#) for the Node Type.
- 1773 • **requirement_definitions**: represents the optional [sequenced](#) list of [requirement definitions](#) for
- 1774 the Node Type.
- 1775 • **capability_definitions**: represents the optional list of [capability definitions](#) for the Node
- 1776 Type.
- 1777 • **interface_definitions**: represents the optional list of one or more [interface definitions](#)
- 1778 supported by the Node Type.
- 1779 • **artifact_definitions**: represents the optional list of [artifact definitions](#) for the Node Type.

1780 3.6.9.3 Additional Requirements

- 1781 • Requirements are intentionally expressed as a sequenced list of TOSCA [Requirement definitions](#)
- 1782 which **SHOULD** be resolved (processed) in sequence order by TOSCA Orchestrators. .

1783 3.6.9.4 Best Practices

- 1784 • It is recommended that all Node Types **SHOULD** derive directly (as a parent) or indirectly (as an
- 1785 ancestor) of the TOSCA Root Node Type (i.e., `tosca.nodes.Root`) to promote compatibility and
- 1786 portability. However, it is permitted to author Node Types that do not do so.
- 1787 • TOSCA Orchestrators, having a full view of the complete application topology template and its
- 1788 resultant dependency graph of nodes and relationships, **MAY** prioritize how they instantiate the nodes
- 1789 and relationships for the application (perhaps in parallel where possible) to achieve the greatest
- 1790 efficiency

1791 **3.6.9.5 Example**

```
my_company.my_types.my_app_node_type:
  derived_from: toska.nodes.SoftwareComponent
  description: My company's custom applicaton
  properties:
    my_app_password:
      type: string
      description: application password
      constraints:
        - min_length: 6
        - max_length: 10
  attributes:
    my_app_port:
      type: integer
      description: application port number
  requirements:
    - some_database:
        capability: EndPoint.Database
        node: Database
        relationship: ConnectsTo
```

1792 **3.6.10 Relationship Type**

1793 A Relationship Type is a reusable entity that defines the type of one or more relationships between Node
1794 Types or Node Templates.

1795 **3.6.10.1 Keynames**

1796 The Relationship Type is a TOSCA Entity and has the common keynames listed in section 3.6.1 TOSCA
1797 Entity Schema.

1798 In addition, the Relationship Type has the following recognized keynames:

Keyname	Required	Definition/Type	Description
properties	no	list of property definitions	An optional list of property definitions for the Relationship Type.
attributes	no	list of attribute definitions	An optional list of attribute definitions for the Relationship Type.
interfaces	no	list of interface definitions	An optional list of interface definitions interfaces supported by the Relationship Type.
valid_target_types	no	string[]	An optional list of one or more names of Capability Types that are valid targets for this relationship.

1799 **3.6.10.2 Grammar**

1800 Relationship Types have following grammar:

```
<relationship type name>:
  derived_from: <parent relationship type name>
  version: <version number>
  metadata:
```

```
<map of string>
description: <relationship description>
properties:
  <property definitions>
attributes:
  <attribute definitions>
interfaces:
  <interface definitions>
valid_target_types: [ <capability type names> ]
```

1801 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1802 • **relationship_type_name**: represents the required symbolic name of the Relationship Type
1803 being declared as a [string](#).
- 1804 • **parent_relationship_type_name**: represents the name ([string](#)) of the [Relationship Type](#) this
1805 Relationship Type definition derives from (i.e., its “parent” type).
- 1806 • **relationship_description**: represents the optional [description](#) string for the corresponding
1807 **relationship_type_name**.
- 1808 • **version_number**: represents the optional TOSCA [version](#) number for the Relationship Type.
- 1809 • **property_definitions**: represents the optional list of [property definitions](#) for the Relationship
1810 Type.
- 1811 • **attribute_definitions**: represents the optional list of [attribute definitions](#) for the Relationship
1812 Type.
- 1813 • **interface_definitions**: represents the optional list of one or more names of valid [interface](#)
1814 [definitions](#) supported by the Relationship Type.
- 1815 • **capability_type_names**: represents one or more names of valid target types for the
1816 relationship (i.e., [Capability Types](#)).

1817 3.6.10.3 Best Practices

- 1818 • For TOSCA application portability, it is recommended that designers use the normative
1819 Relationship types defined in this specification where possible and derive from them for
1820 customization purposes.
- 1821 • The TOSCA Root Relationship Type (**tosca.relationships.Root**) SHOULD be used to derive
1822 new types where possible when defining new relationships types. This assures that its normative
1823 configuration interface (**tosca.interfaces.relationship.Configure**) can be used in a
1824 deterministic way by TOSCA orchestrators.

1825 3.6.10.4 Examples

```
mycompanytypes.myrelationships.AppDependency:
  derived_from: toska.relationships.DependsOn
  valid_target_types: [ mycompanytypes.mycapabilities.SomeAppCapability ]
```

1826 3.6.11 Group Type

1827 A Group Type defines logical grouping types for nodes, typically for different management purposes.
1828 Groups can effectively be viewed as logical nodes that are not part of the physical deployment topology of
1829 an application, yet can have capabilities and the ability to attach policies and interfaces that can be
1830 applied (depending on the group type) to its member nodes.

1831

1832 Conceptually, group definitions allow the creation of logical “membership” relationships to nodes in a
 1833 service template that are not a part of the application’s explicit requirement dependencies in the topology
 1834 template (i.e. those required to actually get the application deployed and running). Instead, such logical
 1835 membership allows for the introduction of things such as group management and uniform application of
 1836 policies (i.e., requirements that are also not bound to the application itself) to the group’s members.

1837 3.6.11.1 Keynames

1838 The Group Type is a TOSCA Entity and has the common keynames listed in section 3.6.1 TOSCA Entity
 1839 Schema.

1840 In addition, the Group Type has the following recognized keynames:

Keyname	Required	Type	Description
attributes	no	list of attribute definitions	An optional list of attribute definitions for the Group Type.
properties	no	list of property definitions	An optional list of property definitions for the Group Type.
members	no	string []	An optional list of one or more names of Node Types that are valid (allowed) as members of the Group Type. Note: This can be viewed by TOSCA Orchestrators as an implied relationship from the listed members nodes to the group, but one that does not have operational lifecycle considerations. For example, if we were to name this as an explicit Relationship Type we might call this “MemberOf” (group).
requirements	no	list of requirement definitions	An optional <i>sequenced</i> list of requirement definitions for the Group Type.
capabilities	no	list of capability definitions	An optional list of capability definitions for the Group Type.
interfaces	no	list of interface definitions	An optional list of interface definitions supported by the Group Type.

1841 3.6.11.2 Grammar

1842 Group Types have one the following grammars:

```

<group type name>:
  derived_from: <parent group type name>
  version: <version number>
  metadata:
    <map of string>
  description: <group description>
  properties:
    <property definitions>
  members: [ <list_of_valid_member_types> ]
  requirements:
    - <requirement definitions>
  capabilities:
    <capability definitions>
  interfaces:
    <interface definitions>
  
```

1843 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1844 • **group_type_name**: represents the required symbolic name of the Group Type being declared as
1845 a [string](#).
- 1846 • **parent_group_type_name**: represents the name ([string](#)) of the [Group Type](#) this Group Type
1847 definition derives from (i.e., its “parent” type).
- 1848 • **version_number**: represents the optional TOSCA [version](#) number for the Group Type.
- 1849 • **group_description**: represents the optional description string for the corresponding
1850 **group_type_name**.
- 1851 • **property_definitions**: represents the optional list of [property definitions](#) for the Group Type.
- 1852 • **list_of_valid_member_types**: represents the optional list of TOSCA types (e.g., Node,
1853 Capability or even other Group Types) that are valid member types for being added to (i.e.,
1854 members of) the Group Type.
- 1855 • **interface_definitions**: represents the optional list of one or more [interface definitions](#)
1856 supported by the Group Type.

1857 3.6.11.3 Additional Requirements

- 1858 • Group definitions **SHOULD NOT** be used to define or redefine relationships (dependencies) for
1859 an application that can be expressed using normative TOSCA Relationships within a TOSCA
1860 topology template.
- 1861 • The list of values associated with the “members” keyname **MUST** only contain types that or
1862 homogenous (i.e., derive from the same type hierarchy).

1863 3.6.11.4 Example

1864 The following represents a Group Type definition:

```
group_types:
  mycompany.mytypes.groups.placement:
    description: My company's group type for placing nodes of type Compute
    members: [ tosca.nodes.Compute ]
```

1865 3.6.12 Policy Type

1866 A Policy Type defines a type of requirement that affects or governs an application or service's topology at
1867 some stage of its lifecycle, but is not explicitly part of the topology itself (i.e., it does not prevent the
1868 application or service from being deployed or run if it did not exist).

1869 3.6.12.1 Keynames

1870 The Policy Type is a TOSCA Entity and has the common keynames listed in section 3.6.1 TOSCA Entity
1871 Schema.

1872 In addition, the Policy Type has the following recognized keynames:

Keyname	Required	Type	Description
properties	no	list of property definitions	An optional list of property definitions for the Policy Type.
targets	no	string []	An optional list of valid Node Types or Group Types the Policy Type can be applied to. Note: This can be viewed by TOSCA Orchestrators as an implied relationship to the target nodes, but one that does not have operational lifecycle considerations. For example, if we were to name this as an explicit Relationship Type we might call this “AppliesTo” (node or group).

Keyname	Required	Type	Description
triggers	no	list of trigger	An optional list of policy triggers for the Policy Type.

1873 3.6.12.2 Grammar

1874 Policy Types have the following grammar:

```
<policy_type_name>:
  derived_from: <parent_policy_type_name>
  version: <version_number>
  metadata:
    <map of string>
  description: <policy_description>
  properties:
    <property_definitions>
  targets: [ <list_of_valid_target_types> ]
  triggers:
    <list_of_trigger_definitions>
```

1875 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1876 • **policy_type_name**: represents the required symbolic name of the Policy Type being declared
- 1877 as a [string](#).
- 1878 • **parent_policy_type_name**: represents the name ([string](#)) of the Policy Type this Policy Type
- 1879 definition derives from (i.e., its “parent” type).
- 1880 • **version_number**: represents the optional TOSCA [version](#) number for the Policy Type.
- 1881 • **policy_description**: represents the optional description string for the corresponding
- 1882 **policy_type_name**.
- 1883 • **property_definitions**: represents the optional list of [property definitions](#) for the Policy Type.
- 1884 • **list_of_valid_target_types**: represents the optional list of TOSCA types (i.e., Group or
- 1885 Node Types) that are valid targets for this Policy Type.
- 1886 • **list_of_trigger_definitions**: represents the optional list of [trigger definitions](#) for the policy.

1887 3.6.12.3 Example

1888 The following represents a Policy Type definition:

```
policy_types:
  mycompany.mytypes.policies.placement.Container.Linux:
    description: My company's placement policy for linux
    derived_from: tosa.policies.Root
```

1889 3.7 Template-specific definitions

1890 The definitions in this section provide reusable modeling element grammars that are specific to the Node
1891 or Relationship templates.

1892 3.7.1 Capability assignment

1893 A capability assignment allows node template authors to assign values to properties and attributes for a
1894 named capability definition that is part of a Node Template’s type definition.

1895 3.7.1.1 Keynames

1896 The following is the list of recognized keynames for a TOSCA capability assignment:

Keyname	Required	Type	Description
properties	no	list of property assignments	An optional list of property definitions for the Capability definition.
attributes	no	list of attribute assignments	An optional list of attribute definitions for the Capability definition.

1897 3.7.1.2 Grammar

1898 Capability assignments have one of the following grammars:

```
<capability_definition_name>:
  properties:
    <property_assignments>
  attributes:
    <attribute_assignments>
```

1899 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1900 • **capability_definition_name**: represents the symbolic name of the capability as a [string](#).
- 1901 • **property_assignments**: represents the optional list of [property assignments](#) for the capability definition.
- 1902 • **attribute_assignments**: represents the optional list of [attribute assignments](#) for the capability definition.

1905 3.7.1.3 Example

1906 The following example shows a capability assignment:

1907 3.7.1.3.1 Notation example

```
node_templates:
  some_node_template:
    capabilities:
      some_capability:
        properties:
          limit: 100
```

1908 3.7.2 Requirement assignment

1909 A Requirement assignment allows template authors to provide either concrete names of TOSCA
1910 templates or provide abstract selection criteria for providers to use to find matching TOSCA templates
1911 that are used to fulfill a named requirement's declared TOSCA Node Type.

1912 3.7.2.1 Keynames

1913 The following is the list of recognized keynames for a TOSCA requirement assignment:

Keyname	Required	Type	Description
capability	no	string	The optional reserved keyname used to provide the name of either a: <ul style="list-style-type: none"> • Capability definition within a <i>target</i> node template that can fulfill the requirement. • Capability Type that the provider will use to select a type-compatible <i>target</i> node template to fulfill the requirement at runtime.
node	no	string	The optional reserved keyname used to identify the target node of a relationship. specifically, it is used to provide either a: <ul style="list-style-type: none"> • Node Template name that can fulfill the target node requirement. • Node Type name that the provider will use to select a type-compatible node template to fulfill the requirement at runtime.
relationship	no	string	The optional reserved keyname used to provide the name of either a: <ul style="list-style-type: none"> • Relationship Template to use to relate the <i>source</i> node to the (capability in the) <i>target</i> node when fulfilling the requirement. • Relationship Type that the provider will use to select a type-compatible relationship template to relate the <i>source</i> node to the <i>target</i> node at runtime.
node_filter	no	node filter	The optional filter definition that TOSCA orchestrators or providers would use to select a type-compatible <i>target</i> node that can fulfill the associated abstract requirement at runtime.

1914 The following is the list of recognized keynames for a TOSCA requirement assignment's **relationship**
1915 keyname which is used when Property assignments need to be provided to inputs of declared interfaces
1916 or their operations:

Keyname	Required	Type	Description
type	no	string	The optional reserved keyname used to provide the name of the Relationship Type for the requirement assignment's relationship keyname.
properties	no	list of interface definitions	The optional reserved keyname used to reference declared (named) interface definitions of the corresponding Relationship Type in order to provide Property assignments for these interfaces or operations of these interfaces.

1917 3.7.2.2 Grammar

1918 Named requirement assignments have one of the following grammars:

1919 3.7.2.2.1 Short notation:

1920 The following single-line grammar may be used if only a concrete Node Template for the target node
1921 needs to be declared in the requirement:

```
<requirement name>: <node template name>
```

1922 This notation is only valid if the corresponding Requirement definition in the Node Template's parent
1923 Node Type declares (at a minimum) a valid Capability Type which can be found in the declared target
1924 Node Template. A valid capability definition always needs to be provided in the requirement declaration of
1925 the *source* node to identify a specific capability definition in the *target* node the requirement will form a
1926 TOSCA relationship with.

1927 **3.7.2.2.2 Extended notation:**

1928 The following grammar would be used if the requirement assignment needs to provide more information
1929 than just the Node Template name:

```
<requirement_name>:  
  node: <node_template_name> | <node_type_name>  
  relationship: <relationship_template_name> | <relationship_type_name>  
  capability: <capability_symbolic_name> | <capability_type_name>  
  node_filter:  
    <node_filter_definition>  
  occurrences: [ min_occurrences, max_occurrences ]
```

1930 **3.7.2.2.3 Extended grammar with Property Assignments for the relationship's**
1931 **Interfaces**

1932 The following additional multi-line grammar is provided for the relationship keyname in order to provide
1933 new Property assignments for inputs of known Interface definitions of the declared Relationship Type.

```
<requirement_name>:  
  # Other keynames omitted for brevity  
  relationship:  
    type: <relationship_template_name> | <relationship_type_name>  
    properties:  
      <property_assignments>  
    interfaces:  
      <interface_assignments>
```

1934 Examples of uses for the extended requirement assignment grammar include:

- 1935 • The need to allow runtime selection of the target node based upon an abstract Node Type rather
1936 than a concrete Node Template. This may include use of the node_filter keyname to provide
1937 node and capability filtering information to find the “best match” of a concrete Node Template at
1938 runtime.
- 1939 • The need to further clarify the concrete Relationship Template or abstract Relationship Type to
1940 use when relating the source node's requirement to the target node's capability.
- 1941 • The need to further clarify the concrete capability (symbolic) name or abstract Capability Type in
1942 the target node to form a relationship between.
- 1943 • The need to (further) constrain the occurrences of the requirement in the instance model.

1944 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1945 • **requirement_name**: represents the symbolic name of a requirement assignment as a [string](#).
- 1946 • **node_template_name**: represents the optional name of a Node Template that contains the
1947 capability this requirement will be fulfilled by.
- 1948 • **relationship_template_name**: represents the optional name of a [Relationship Type](#) to be used
1949 when relating the requirement appears to the capability in the target node.
- 1950 • **capability_symbolic_name**: represents the optional ordered list of specific, required capability
1951 type or named capability definition within the target Node Type or Template.
- 1952 • **node_type_name**: represents the optional name of a TOSCA Node Type the associated named
1953 requirement can be fulfilled by. This must be a type that is compatible with the Node Type
1954 declared on the matching requirement (same symbolic name) the requirement's Node Template
1955 is based upon.
- 1956 • **relationship_type_name**: represents the optional name of a [Relationship Type](#) that is
1957 compatible with the Capability Type in the target node.

- 1958 • **property_assignments**: represents the optional list of property value assignments for the
1959 declared relationship.
- 1960 • **interface_assignments**: represents the optional list of interface definitions for the declared
1961 relationship used to provide property assignments on inputs of interfaces and operations.
- 1962 • **capability_type_name**: represents the optional name of a Capability Type definition within the
1963 target Node Type this requirement needs to form a relationship with.
- 1964 • **node_filter_definition**: represents the optional [node filter](#) TOSCA orchestrators would use
1965 to fulfill the requirement for selecting a target node. Note that this SHALL only be valid if the **node**
1966 keyname's value is a Node Type and is invalid if it is a Node Template.

1967 3.7.2.3 Examples

1968 3.7.2.3.1 Example 1 – Abstract hosting requirement on a Node Type

1969 A web application node template named 'my_application_node_template' of type **WebApplication**
1970 declares a requirement named 'host' that needs to be fulfilled by any node that derives from the node
1971 type **WebServer**.

```
# Example of a requirement fulfilled by a specific web server node template
node_templates:
  my_application_node_template:
    type: tosca.nodes.WebApplication
    ...
  requirements:
    - host:
      node: tosca.nodes.WebServer
```

1972 In this case, the node template's type is **WebApplication** which already declares the Relationship Type
1973 **HostedOn** to use to relate to the target node and the Capability Type of **Container** to be the specific
1974 target of the requirement in the target node.

1975 3.7.2.3.2 Example 2 - Requirement with Node Template and a custom Relationship 1976 Type

1977 This example is similar to the previous example; however, the requirement named 'database' describes
1978 a requirement for a connection to a database endpoint (**Endpoint.Database**) Capability Type in a named
1979 node template (**my_database**). However, the connection requires a custom Relationship Type
1980 (**my.types.CustomDbConnection**) declared on the keyname 'relationship'.

```
# Example of a (database) requirement that is fulfilled by a node template named
# "my_database", but also requires a custom database connection relationship
my_application_node_template:
  requirements:
    - database:
      node: my_database
      capability: Endpoint.Database
      relationship: my.types.CustomDbConnection
```

1981 3.7.2.3.3 Example 3 - Requirement for a Compute node with additional selection 1982 criteria (filter)

1983 This example shows how to extend an abstract 'host' requirement for a **Compute** node
1984 with a filter definition that further constrains TOSCA orchestrators to include
1985 additional properties and capabilities on the target node when fulfilling the
1986 requirement.

```

node_templates:
  mysql:
    type: tosca.nodes.DBMS.MySQL
    properties:
      # omitted here for brevity
    requirements:
      - host:
          node: tosca.nodes.Compute
          node_filter:
            capabilities:
              - host:
                  properties:
                    - num_cpus: { in_range: [ 1, 4 ] }
                    - mem_size: { greater_or_equal: 512 MB }
              - os:
                  properties:
                    - architecture: { equal: x86_64 }
                    - type: { equal: linux }
                    - distribution: { equal: ubuntu }
            - mytypes.capabilities.compute.encryption:
                  properties:
                    - algorithm: { equal: aes }
                    - keylength: { valid_values: [ 128, 256 ] }

```

1987 **3.7.3 Node Template**

1988 A Node Template specifies the occurrence of a manageable software component as part of an
 1989 application's topology model which is defined in a TOSCA Service Template. A Node template is an
 1990 instance of a specified Node Type and can provide customized properties, constraints or operations
 1991 which override the defaults provided by its Node Type and its implementations.

1992 **3.7.3.1 Keynames**

1993 The following is the list of recognized keynames for a TOSCA Node Template definition:

Keyname	Required	Type	Description
type	yes	string	The required name of the Node Type the Node Template is based upon.
description	no	description	An optional description for the Node Template.
metadata	no	map of string	Defines a section used to declare additional metadata information.
directives	no	string[]	An optional list of directive values to provide processing instructions to orchestrators and tooling.
properties	no	list of property assignments	An optional list of property value assignments for the Node Template.
attributes	no	list of attribute assignments	An optional list of attribute value assignments for the Node Template.
requirements	no	list of requirement assignments	An optional <u>sequenced</u> list of requirement assignments for the Node Template.

Keyname	Required	Type	Description
capabilities	no	list of capability assignments	An optional list of capability assignments for the Node Template.
interfaces	no	list of interface definitions	An optional list of named interface definitions for the Node Template.
artifacts	no	list of artifact definitions	An optional list of named artifact definitions for the Node Template.
node_filter	no	node filter	The optional filter definition that TOSCA orchestrators would use to select the correct target node. This keyname is only valid if the directive has the value of “selectable” set.
copy	no	string	The optional (symbolic) name of another node template to copy into (all keynames and values) and use as a basis for this node template.

1994 3.7.3.2 Grammar

```

<node_template_name>:
  type: <node_type_name>
  description: <node_template_description>
  directives: [<directives>]
  metadata:
    <map of string>
  properties:
    <property_assignments>
  attributes:
    <attribute_assignments>
  requirements:
    - <requirement_assignments>
  capabilities:
    <capability_assignments>
  interfaces:
    <interface_definitions>
  artifacts:
    <artifact_definitions>
  node_filter:
    <node_filter_definition>
  copy: <source_node_template_name>

```

1995 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1996 • **node_template_name**: represents the required symbolic name of the Node Template being
- 1997 declared.
- 1998 • **node_type_name**: represents the name of the Node Type the Node Template is based upon.
- 1999 • **node_template_description**: represents the optional [description](#) string for Node Template.
- 2000 • **directives**: represents the optional list of processing instruction keywords (as strings) for use by
- 2001 tooling and orchestrators.
- 2002 • **property_assignments**: represents the optional list of [property assignments](#) for the Node
- 2003 Template that provide values for properties defined in its declared Node Type.
- 2004 • **attribute_assignments**: represents the optional list of [attribute assignments](#) for the Node
- 2005 Template that provide values for attributes defined in its declared Node Type.

- 2006 • **requirement_assignments**: represents the optional *sequenced* list of [requirement assignments](#)
- 2007 for the Node Template that allow assignment of type-compatible capabilities, target nodes,
- 2008 relationships and target (node filters) for use when fulfilling the requirement at runtime.
- 2009 • **capability_assignments**: represents the optional list of [capability assignments](#) for the Node
- 2010 Template that augment those provided by its declared Node Type.
- 2011 • **interface_definitions**: represents the optional list of [interface definitions](#) for the Node
- 2012 Template that *augment* those provided by its declared Node Type.
- 2013 • **artifact_definitions**: represents the optional list of [artifact definitions](#) for the Node Template
- 2014 that augment those provided by its declared Node Type.
- 2015 • **node_filter_definition**: represents the optional [node filter](#) TOSCA orchestrators would use
- 2016 for selecting a matching node template.
- 2017 • **source_node_template_name**: represents the optional (symbolic) name of another node
- 2018 template to copy into (all keynames and values) and use as a basis for this node template.

2019 3.7.3.3 Additional requirements

- 2020 • The **node_filter** keyword (and supporting grammar) **SHALL** only be valid if the Node Template
- 2021 has a **directive** keyname with the value of “**selectable**” set.
- 2022 • The source node template provided as a value on the **copy** keyname **MUST NOT** itself use the
- 2023 **copy** keyname (i.e., it must itself be a complete node template description and not copied from
- 2024 another node template).

2025 3.7.3.4 Example

```
node_templates:
  mysql:
    type: tosca.nodes.DBMS.MySQL
    properties:
      root_password: { get_input: my_mysql_rootpw }
      port: { get_input: my_mysql_port }
    requirements:
      - host: db_server
    interfaces:
      Standard:
        configure: scripts/my_own_configure.sh
```

2026 3.7.4 Relationship Template

2027 A Relationship Template specifies the occurrence of a manageable relationship between node templates
 2028 as part of an application’s topology model that is defined in a TOSCA Service Template. A Relationship
 2029 template is an instance of a specified Relationship Type and can provide customized properties,
 2030 constraints or operations which override the defaults provided by its Relationship Type and its
 2031 implementations.

2032 3.7.4.1 Keynames

2033 The following is the list of recognized keynames for a TOSCA Relationship Template definition:

Keyname	Required	Type	Description
type	yes	string	The required name of the Relationship Type the Relationship Template is based upon.
description	no	description	An optional description for the Relationship Template.

Keyname	Required	Type	Description
metadata	no	map of string	Defines a section used to declare additional metadata information.
properties	no	list of property assignments	An optional list of property assignments for the Relationship Template.
attributes	no	list of attribute assignments	An optional list of attribute assignments for the Relationship Template.
interfaces	no	list of interface definitions	An optional list of named interface definitions for the Node Template.
copy	no	string	The optional (symbolic) name of another relationship template to copy into (all keynames and values) and use as a basis for this relationship template.

2034 3.7.4.2 Grammar

```

<relationship_template_name>:
  type: <relationship type name>
  description: <relationship type description>
  metadata:
    <map of string>
  properties:
    <property assignments>
  attributes:
    <attribute assignments>
  interfaces:
    <interface definitions>
  copy:
    <source relationship template name>

```

2035 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 2036 • **relationship_template_name**: represents the required symbolic name of the Relationship
- 2037 Template being declared.
- 2038 • **relationship_type_name**: represents the name of the Relationship Type the Relationship
- 2039 Template is based upon.
- 2040 • **relationship_template_description**: represents the optional [description](#) string for the
- 2041 Relationship Template.
- 2042 • **property_assignments**: represents the optional list of [property assignments](#) for the Relationship
- 2043 Template that provide values for properties defined in its declared Relationship Type.
- 2044 • **attribute_assignments**: represents the optional list of [attribute assignments](#) for the
- 2045 Relationship Template that provide values for attributes defined in its declared Relationship Type.
- 2046 • **interface_definitions**: represents the optional list of [interface definitions](#) for the Relationship
- 2047 Template that augment those provided by its declared Relationship Type.
- 2048 • **source_relationship_template_name**: represents the optional (symbolic) name of another
- 2049 relationship template to copy into (all keynames and values) and use as a basis for this
- 2050 relationship template.

2051 **3.7.4.3 Additional requirements**

- 2052
- The source relationship template provided as a value on the **copy** keyname MUST NOT itself use the **copy** keyname (i.e., it must itself be a complete relationship template description and not copied from another relationship template).
- 2053
- 2054

2055 **3.7.4.4 Example**

```
relationship_templates:  
  storage_attachment:  
    type: AttachesTo  
    properties:  
      location: /my_mount_point
```

2056 **3.7.5 Group definition**

2057 A group definition defines a logical grouping of node templates, typically for management purposes, but is
2058 separate from the application's topology template.

2059 **3.7.5.1 Keynames**

2060 The following is the list of recognized keynames for a TOSCA group definition:

Keyname	Required	Type	Description
type	yes	string	The required name of the group type the group definition is based upon.
description	no	description	The optional description for the group definition.
metadata	no	map of string	Defines a section used to declare additional metadata information.
properties	no	list of property assignments	An optional list of property value assignments for the group definition.
members	no	list of string	The optional list of one or more node template names that are members of this group definition.
interfaces	no	list of interface definitions	An optional list of named interface definitions for the group definition.

2061 **3.7.5.2 Grammar**

2062 Group definitions have one the following grammars:

```
<group\_name>:  
  type: <group\_type\_name>  
  description: <group\_description>  
  metadata:  
    <map of string>  
  properties:  
    <property assignments>  
  members: [ <list\_of\_node\_templates> ]  
  interfaces:  
    <interface definitions>
```

2063 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 2064
- **group_name**: represents the required symbolic name of the group as a [string](#).

- 2065 • **group_type_name**: represents the name of the Group Type the definition is based upon.
- 2066 • **group_description**: contains an optional description of the group.
- 2067 • **property_assignments**: represents the optional list of [property assignments](#) for the group
- 2068 definition that provide values for properties defined in its declared Group Type.
- 2069 • **list_of_node_templates**: contains the required list of one or more node template names
- 2070 (within the same topology template) that are members of this logical group.
- 2071 • **interface_definitions**: represents the optional list of [interface definitions](#) for the group
- 2072 definition that augment those provided by its declared Group Type.

2073 3.7.5.3 Additional Requirements

- 2074 • Group definitions **SHOULD NOT** be used to define or redefine relationships (dependencies) for
- 2075 an application that can be expressed using normative TOSCA Relationships within a TOSCA
- 2076 topology template.

2077 3.7.5.4 Example

2078 The following represents a group definition:

```
groups:
  my_app_placement_group:
    type: tosca.groups.Root
    description: My application's logical component grouping for placement
    members: [ my_web_server, my_sql_database ]
```

2079 3.7.6 Policy definition

2080 A policy definition defines a policy that can be associated with a TOSCA topology or top-level entity

2081 definition (e.g., group definition, node template, etc.).

2082 3.7.6.1 Keynames

2083 The following is the list of recognized keynames for a TOSCA policy definition:

Keyname	Required	Type	Description
type	yes	string	The required name of the policy type the policy definition is based upon.
description	no	description	The optional description for the policy definition.
metadata	no	map of string	Defines a section used to declare additional metadata information.
properties	no	list of property assignments	An optional list of property value assignments for the policy definition.
targets	no	string[]	An optional list of valid Node Templates or Groups the Policy can be applied to.

2084 3.7.6.2 Grammar

2085 Policy definitions have one the following grammars:

```
<policy_name>:
  type: <policy_type_name>
  description: <policy_description>
```

```

metadata:
  <map of string>
properties:
  <property assignments>
targets: [<list_of_policy_targets>]
triggers:
  <list_of_trigger_definitions>

```

2086 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 2087 • **policy_name**: represents the required symbolic name of the policy as a [string](#).
- 2088 • **policy_type_name**: represents the name of the policy the definition is based upon.
- 2089 • **policy_description**: contains an optional description of the policy.
- 2090 • **property_assignments**: represents the optional list of [property assignments](#) for the policy
- 2091 definition that provide values for properties defined in its declared Policy Type.
- 2092 • **list_of_policy_targets**: represents the optional list of names of node templates or groups
- 2093 that the policy is to applied to.
- 2094 • **list_of_trigger_definitions**: represents the optional list of [trigger definitions](#) for the policy.

2095 3.7.6.3 Example

2096 The following represents a policy definition:

```

policies:
- my_compute_placement_policy:
  type: toscapolicies.placement
  description: Apply my placement policy to my application's servers
  targets: [ my_server_1, my_server_2 ]
  # remainder of policy definition left off for brevity

```

2097 3.7.7 Imperative Workflow definition

2098 A workflow definition defines an imperative workflow that is associated with a TOSCA topology.

2099 3.7.7.1 Keynames

2100 The following is the list of recognized keynames for a TOSCA workflow definition:

Keyname	Required	Type	Description
description	no	description	The optional description for the workflow definition.
metadata	no	map of string	Defines a section used to declare additional metadata information.
inputs	no	list of property definitions	The optional list of input parameter definitions.
preconditions	no	list of precondition definitions	List of preconditions to be validated before the workflow can be processed.
steps	No	list of step definitions	An optional list of valid Node Templates or Groups the Policy can be applied to.

2101

2102 3.7.7.2 Grammar

2103 Imperative workflow definitions have the following grammar:

```

<workflow_name>:
  description: <workflow_description>
  metadata:
    <map of string>
  inputs:
    <property_definitions>
  preconditions:
    - <workflow_precondition_definition>
  steps:
    <workflow_steps>

```

2104 In the above grammar, the pseudo values that appear in angle

2105 3.8 Topology Template definition

2106 This section defines the topology template of a cloud application. The main ingredients of the topology
 2107 template are node templates representing components of the application and relationship templates
 2108 representing links between the components. These elements are defined in the nested **node_templates**
 2109 section and the nested **relationship_templates** sections, respectively. Furthermore, a topology
 2110 template allows for defining input parameters, output parameters as well as grouping of node templates.

2111 3.8.1 Keynames

2112 The following is the list of recognized keynames for a TOSCA Topology Template:

Keyname	Required	Type	Description
description	no	description	The optional description for the Topology Template.
inputs	no	list of parameter definitions	An optional list of input parameters (i.e., as parameter definitions) for the Topology Template.
node_templates	no	list of node templates	An optional list of node template definitions for the Topology Template.
relationship_templates	no	list of relationship templates	An optional list of relationship templates for the Topology Template.
groups	no	list of group definitions	An optional list of Group definitions whose members are node templates defined within this same Topology Template.
policies	no	list of policy definitions	An optional list of Policy definitions for the Topology Template.
outputs	no	list of parameter definitions	An optional list of output parameters (i.e., as parameter definitions) for the Topology Template.

Keyname	Required	Type	Description
substitution_mappings	no	N/A	An optional declaration that exports the topology template as an implementation of a Node type. This also includes the mappings between the external Node Types named capabilities and requirements to existing implementations of those capabilities and requirements on Node templates declared within the topology template.
workflows	no	list of imperative workflow definitions	An optional map of imperative workflow definition for the Topology Template.

2113 3.8.2 Grammar

2114 The overall grammar of the **topology_template** section is shown below.—Detailed grammar definitions
2115 of the each sub-sections are provided in subsequent subsections.

```

topology_template:
  description: <template_description>
  inputs: <input_parameter_list>
  outputs: <output_parameter_list>
  node_templates: <node_template_list>
  relationship_templates: <relationship_template_list>
  groups: <group_definition_list>
  policies:
    - <policy_definition_list>
  workflows: <workflow_list>
  # Optional declaration that exports the Topology Template
  # as an implementation of a Node Type.
  substitution_mappings:
    node_type: <node_type_name>
    capabilities:
      <map_of_capability_mappings_to_expose>
    requirements:
      <map_of_requirement_mapping_to_expose>

```

2116 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 2117 • **template_description**: represents the optional [description](#) string for Topology Template.
- 2118 • **input_parameter_list**: represents the optional list of input parameters (i.e., as property
2119 definitions) for the Topology Template.
- 2120 • **output_parameter_list**: represents the optional list of output parameters (i.e., as property
2121 definitions) for the Topology Template.
- 2122 • **group_definition_list**: represents the optional list of [group definitions](#) whose members are
2123 node templates that also are defined within this Topology Template.
- 2124 • **policy_definition_list**: represents the optional sequenced list of policy definitions for the
2125 Topology Template.
- 2126 • **workflow_list**: represents the optional list of imperative workflow definitions
2127 for the Topology Template.

- 2128 • **node_template_list**: represents the optional list of [node template](#) definitions for the Topology
2129 Template.
- 2130 • **relationship_template_list**: represents the optional list of [relationship templates](#) for the
2131 Topology Template.
- 2132 • **node_type_name**: represents the optional name of a [Node Type](#) that the Topology Template
2133 implements as part of the **substitution_mappings**.
- 2134 • **map_of_capability_mappings_to_expose**: represents the mappings that expose internal
2135 capabilities from node templates (within the topology template) as capabilities of the Node Type
2136 definition that is declared as part of the **substitution_mappings**.
- 2137 • **map_of_requirement_mappings_to_expose**: represents the mappings of link requirements of
2138 the Node Type definition that is declared as part of the **substitution_mappings** to internal
2139 requirements implementations within node templates (declared within the topology template).
2140

2141 More detailed explanations for each of the Topology Template grammar's keynames appears in the
2142 sections below.

2143 3.8.2.1 inputs

2144 The **inputs** section provides a means to define parameters using TOSCA parameter definitions, their
2145 allowed values via constraints and default values within a TOSCA Simple Profile template. Input
2146 parameters defined in the **inputs** section of a topology template can be mapped to properties of node
2147 templates or relationship templates within the same topology template and can thus be used for
2148 parameterizing the instantiation of the topology template.

2149

2150 This section defines topology template-level input parameter section.

- 2151 • Inputs here would ideally be mapped to BoundaryDefinitions in TOSCA v1.0.
- 2152 • Treat input parameters as fixed global variables (not settable within template)
- 2153 • If not in input take default (nodes use default)

2154 3.8.2.1.1 Grammar

2155 The grammar of the **inputs** section is as follows:

```
inputs:
  <parameter_definition_list>
```

2156 3.8.2.1.2 Examples

2157 This section provides a set of examples for the single elements of a topology template.

2158 Simple **inputs** example without any constraints:

```
inputs:
  fooName:
    type: string
    description: Simple string typed property definition with no constraints.
    default: bar
```

2159 Example of **inputs** with constraints:

```
inputs:
  SiteName:
    type: string
    description: string typed property definition with constraints
```



```
default: My Site
constraints:
  - min_length: 9
```

2160 **3.8.2.2 node_templates**

2161 The `node_templates` section lists the Node Templates that describe the (software) components that are
2162 used to compose cloud applications.

2163 **3.8.2.2.1 grammar**

2164 The grammar of the `node_templates` section is as follows:

```
node_templates:
  <node_template_defn_1>
  ...
  <node_template_defn_n>
```

2165 **3.8.2.2.2 Example**

2166 Example of `node_templates` section:

```
node_templates:
  my_webapp_node_template:
    type: WebApplication

  my_database_node_template:
    type: Database
```

2167 **3.8.2.3 relationship_templates**

2168 The `relationship_templates` section lists the Relationship Templates that describe the relations
2169 between components that are used to compose cloud applications.

2170

2171 Note that in the TOSCA Simple Profile, the explicit definition of relationship templates as it was required
2172 in TOSCA v1.0 is optional, since relationships between nodes get implicitly defined by referencing other
2173 node templates in the requirements sections of node templates.

2174 **3.8.2.3.1 Grammar**

2175 The grammar of the `relationship_templates` section is as follows:

```
relationship_templates:
  <relationship_template_defn_1>
  ...
  <relationship_template_defn_n>
```

2176 **3.8.2.3.2 Example**

2177 Example of `relationship_templates` section:

```
relationship_templates:
  my_connectsto_relationship:
    type: tosca.relationships.ConnectsTo
    interfaces:
      Configure:
        inputs:
```

```
speed: { get_attribute: [ SOURCE, connect_speed ] }
```

2178 3.8.2.4 outputs

2179 The **outputs** section provides a means to define the output parameters that are available from a TOSCA
2180 Simple Profile service template. It allows for exposing attributes of node templates or relationship
2181 templates within the containing **topology_template** to users of a service.

2182 3.8.2.4.1 Grammar

2183 The grammar of the **outputs** section is as follows:

```
outputs:  
  <parameter_def_list>
```

2184 3.8.2.4.2 Example

2185 Example of the **outputs** section:

```
outputs:  
  server_address:  
    description: The first private IP address for the provisioned server.  
    value: { get_attribute: [ HOST, networks, private, addresses, 0 ] }
```

2186 3.8.2.5 groups

2187 The **groups** section allows for grouping one or more node templates within a TOSCA Service Template
2188 and for assigning special attributes like policies to the group.

2189 3.8.2.5.1 Grammar

2190 The grammar of the **groups** section is as follows:

```
groups:  
  <group_defn 1>  
  ...  
  <group_defn n>
```

2191 3.8.2.5.2 Example

2192 The following example shows the definition of three Compute nodes in the **node_templates** section of a
2193 **topology_template** as well as the grouping of two of the Compute nodes in a group **server_group_1**.

```
node_templates:  
  server1:  
    type: toasca.nodes.Compute  
    # more details ...  
  
  server2:  
    type: toasca.nodes.Compute  
    # more details ...  
  
  server3:  
    type: toasca.nodes.Compute  
    # more details ...  
  
groups:
```

```
# server2 and server3 are part of the same group
server_group_1:
  type: toska.groups.Root
  members: [ server2, server3 ]
```

2194 3.8.2.6 policies

2195 The **policies** section allows for declaring policies that can be applied to entities in the topology template.

2196 3.8.2.6.1 Grammar

2197 The grammar of the **policies** section is as follows:

```
policies:
- <policy_defn_1>
- ...
- <policy_defn_n>
```

2198 3.8.2.6.2 Example

2199 The following example shows the definition of a placement policy.

```
policies:
- my_placement_policy:
  type: mycompany.mytypes.policy.placement
```

2200 3.8.2.7 Notes

- 2201 • The parameters (properties) that are listed as part of the **inputs** block can be mapped to
- 2202 **PropertyMappings** provided as part of **BoundaryDefinitions** as described by the TOSCA v1.0
- 2203 specification.
- 2204 • The node templates listed as part of the **node_templates** block can be mapped to the list of
- 2205 **NodeTemplate** definitions provided as part of **TopologyTemplate** of a **ServiceTemplate** as
- 2206 described by the TOSCA v1.0 specification.
- 2207 • The relationship templates listed as part of the **relationship_templates** block can be mapped
- 2208 to the list of **RelationshipTemplate** definitions provided as part of **TopologyTemplate** of a
- 2209 **ServiceTemplate** as described by the TOSCA v1.0 specification.
- 2210 • The output parameters that are listed as part of the **outputs** section of a topology template can
- 2211 be mapped to **PropertyMappings** provided as part of **BoundaryDefinitions** as described by
- 2212 the TOSCA v1.0 specification.
- 2213 ○ Note, however, that TOSCA v1.0 does not define a direction (input vs. output) for those
- 2214 mappings, i.e. TOSCA v1.0 **PropertyMappings** are underspecified in that respect and
- 2215 TOSCA Simple Profile's **inputs** and **outputs** provide a more concrete definition of input
- 2216 and output parameters.

2217 3.9 Service Template definition

2218 A TOSCA Service Template (YAML) document contains element definitions of building blocks for cloud

2219 application, or complete models of cloud applications. This section describes the top-level structural

2220 elements (TOSCA keynames) along with their grammars, which are allowed to appear in a TOSCA

2221 Service Template document.

2222 **3.9.1 Keynames**

2223 The following is the list of recognized keynames for a TOSCA Service Template definition:

Keyname	Required	Type	Description
tosca_definitions_version	yes	string	Defines the version of the TOSCA Simple Profile specification the template (grammar) complies with.
metadata	no	map of string	Defines a section used to declare additional metadata information. Domain-specific TOSCA profile specifications may define keynames that are required for their implementations.
description	no	description	Declares a description for this Service Template and its contents.
dsl_definitions	no	N/A	Declares optional DSL-specific definitions and conventions. For example, in YAML, this allows defining reusable YAML macros (i.e., YAML alias anchors) for use throughout the TOSCA Service Template.
repositories	no	list of Repository definitions	Declares the list of external repositories which contain artifacts that are referenced in the service template along with their addresses and necessary credential information used to connect to them in order to retrieve the artifacts.
imports	no	list of Import Definitions	Declares import statements external TOSCA Definitions documents. For example, these may be file location or URIs relative to the service template file within the same TOSCA CSAR file.
artifact_types	no	list of Artifact Types	This section contains an optional list of artifact type definitions for use in the service template
data_types	no	list of Data Types	Declares a list of optional TOSCA Data Type definitions.
capability_types	no	list of Capability Types	This section contains an optional list of capability type definitions for use in the service template.
interface_types	no	list of Interface Types	This section contains an optional list of interface type definitions for use in the service template.
relationship_types	no	list of Relationship Types	This section contains a set of relationship type definitions for use in the service template.
node_types	no	list of Node Types	This section contains a set of node type definitions for use in the service template.
group_types	no	list of Group Types	This section contains a list of group type definitions for use in the service template.
policy_types	no	list of Policy Types	This section contains a list of policy type definitions for use in the service template.
topology_template	no	Topology Template definition	Defines the topology template of an application or service, consisting of node templates that represent the application's or service's components, as well as relationship templates representing relations between the components.

2224 **3.9.1.1 Metadata keynames**

2225 The following is the list of recognized metadata keynames for a TOSCA Service Template definition:

Keyname	Required	Type	Description
template_name	no	string	Declares a descriptive name for the template.
template_author	no	string	Declares the author(s) or owner of the template.
template_version	no	string	Declares the version string for the template.

2226 3.9.2 Grammar

2227 The overall structure of a TOSCA Service Template and its top-level key collations using the TOSCA
2228 Simple Profile is shown below:

```

tosca_definitions_version: # Required TOSCA Definitions version string

# Optional metadata keyname: value pairs
metadata:
  template_name: <value>           # Optional name of this service template
  template_author: <value>         # Optional author of this service template
  template_version: <value>       # Optional version of this service template
  # Optional list of domain or profile specific metadata keynames

# Optional description of the definitions inside the file.
description: <template type description>

dsl_definitions:
  # list of YAML alias anchors (or macros)

repositories:
  # list of external repository definitions which host TOSCA artifacts

imports:
  # ordered list of import definitions

artifact_types:
  # list of artifact type definitions

data_types:
  # list of datatype definitions

capability_types:
  # list of capability type definitions

interface_types
  # list of interface type definitions

relationship_types:
  # list of relationship type definitions

node_types:
  # list of node type definitions

group_types:
  # list of group type definitions

policy_types:
  # list of policy type definitions

```

```
topology_template:
  # topology template definition of the cloud application or service
```

2229 3.9.2.1 Notes

- 2230 • TOSCA Service Templates do not have to contain a topology_template and MAY contain simply
2231 type definitions (e.g., Artifact, Interface, Capability, Node, Relationship Types, etc.) and be
2232 imported for use as type definitions in other TOSCA Service Templates.

2233 3.9.3 Top-level keyname definitions

2234 3.9.3.1 toska_definitions_version

2235 This required element provides a means to include a reference to the TOSCA Simple Profile specification
2236 within the TOSCA Definitions YAML file. It is an indicator for the version of the TOSCA grammar that
2237 should be used to parse the remainder of the document.

2238 3.9.3.1.1 Keyname

```
tosca_definitions_version
```

2239 3.9.3.1.2 Grammar

2240 Single-line form:

```
tosca_definitions_version: <tosca_simple_profile_version>
```

2241 3.9.3.1.3 Examples:

2242 TOSCA Simple Profile version 1.0 specification using the defined namespace alias (see Section 3.1):

```
tosca_definitions_version: toska_simple_yaml_1_0
```

2243 TOSCA Simple Profile version 1.0 specification using the fully defined (target) namespace (see Section
2244 3.1):

```
tosca_definitions_version: http://docs.oasis-open.org/tosca/ns/simple/yaml/1.0
```

2245 3.9.3.2 metadata

2246 This keyname is used to associate domain-specific metadata with the Service Template. The metadata
2247 keyname allows a declaration of a map of keynames with string values.

2248 3.9.3.2.1 Keyname

```
metadata
```

2249 3.9.3.2.2 Grammar

```
metadata:
  <map_of_string_values>
```

2250 3.9.3.2.3 Example

```
metadata:
```

```
creation_date: 2015-04-14
date_updated: 2015-05-01
status: developmental
```

2251

2252 **3.9.3.3 template_name**

2253 This optional metadata keyname can be used to declare the name of service template as a single-line
2254 string value.

2255 **3.9.3.3.1 Keyname**

```
template_name
```

2256 **3.9.3.3.2 Grammar**

```
template_name: <name string>
```

2257 **3.9.3.3.3 Example**

```
template_name: My service template
```

2258 **3.9.3.3.4 Notes**

- 2259
- Some service templates are designed to be referenced and reused by other service templates.
2260 Therefore, in these cases, the **template_name** value SHOULD be designed to be used as a
2261 unique identifier through the use of namespacing techniques.

2262 **3.9.3.4 template_author**

2263 This optional metadata keyname can be used to declare the author(s) of the service template as a single-
2264 line string value.

2265 **3.9.3.4.1 Keyname**

```
template_author
```

2266 **3.9.3.4.2 Grammar**

```
template_author: <author string>
```

2267 **3.9.3.4.3 Example**

```
template_author: My service template
```

2268 **3.9.3.5 template_version**

2269 This optional metadata keyname can be used to declare a domain specific version of the service template
2270 as a single-line string value.

2271 **3.9.3.5.1 Keyname**

```
template_version
```

2272 **3.9.3.5.2 Grammar**

```
template_version: <version>
```

2273 **3.9.3.5.3 Example**

```
template_version: 2.0.17
```

2274 **3.9.3.5.4 Notes:**

- 2275 • Some service templates are designed to be referenced and reused by other service templates
2276 and have a lifecycle of their own. Therefore, in these cases, a **template_version** value
2277 SHOULD be included and used in conjunction with a unique **template_name** value to enable
2278 lifecycle management of the service template and its contents.

2279 **3.9.3.6 description**

2280 This optional keyname provides a means to include single or multiline descriptions within a TOSCA
2281 Simple Profile template as a scalar string value.

2282 **3.9.3.6.1 Keyname**

```
description
```

2283 **3.9.3.7 dsl_definitions**

2284 This optional keyname provides a section to define macros (e.g., YAML-style macros when using the
2285 TOSCA Simple Profile in YAML specification).

2286 **3.9.3.7.1 Keyname**

```
dsl_definitions
```

2287 **3.9.3.7.2 Grammar**

```
dsl_definitions:  
  <dsl\_definition 1>  
  ...  
  <dsl\_definition n>
```

2288 **3.9.3.7.3 Example**

```
dsl_definitions:  
  ubuntu_image_props: &ubuntu_image_props  
    architecture: x86_64  
    type: linux  
    distribution: ubuntu  
    os_version: 14.04  
  
  redhat_image_props: &redhat_image_props  
    architecture: x86_64  
    type: linux  
    distribution: rhel  
    os_version: 6.6
```


2289 3.9.3.8 repositories

2290 This optional keyname provides a section to define external repositories which may contain artifacts or
2291 other TOSCA Service Templates which might be referenced or imported by the TOSCA Service Template
2292 definition.

2293 3.9.3.8.1 Keyname

```
repositories
```

2294 3.9.3.8.2 Grammar

```
repositories:  
  <repository_definition 1>  
  ...  
  <repository_definition n>
```

2295 3.9.3.8.3 Example

```
repositories:  
  my_project_artifact_repo:  
    description: development repository for TAR archives and Bash scripts  
    url: http://mycompany.com/repository/myproject/
```

2296 3.9.3.9 imports

2297 This optional keyname provides a way to import a *block sequence* of one or more TOSCA Definitions
2298 documents. TOSCA Definitions documents can contain reusable TOSCA type definitions (e.g., Node
2299 Types, Relationship Types, Artifact Types, etc.) defined by other authors. This mechanism provides an
2300 effective way for companies and organizations to define normative types and/or describe their software
2301 applications for reuse in other TOSCA Service Templates.

2302 3.9.3.9.1 Keyname

```
imports
```

2303 3.9.3.9.2 Grammar

```
imports:  
  - <import_definition 1>  
  - ...  
  - <import_definition n>
```

2304 3.9.3.9.3 Example

```
# An example import of definitions files from a location relative to the  
# file location of the service template declaring the import.  
imports:  
  - some_definitions: relative_path/my_defns/my_typesdefs_1.yaml  
  - file: my_defns/my_typesdefs_n.yaml  
    repository: my_company_repo  
    namespace_uri: http://mycompany.com/ns/tosca/2.0  
    namespace_prefix: mycompany
```

2305 **3.9.3.10 artifact_types**

2306 This optional keyname lists the Artifact Types that are defined by this Service Template.

2307 **3.9.3.10.1 Keyname**

```
artifact_types
```

2308 **3.9.3.10.2 Grammar**

```
artifact_types:  
  <artifact type defn 1>  
  ...  
  <artifact type defn n>
```

2309 **3.9.3.10.3 Example**

```
artifact_types:  
  mycompany.artifacttypes.myFileType:  
    derived_from: toscartifacts.File
```

2310 **3.9.3.11 data_types**

2311 This optional keyname provides a section to define new data types in TOSCA.

2312 **3.9.3.11.1 Keyname**

```
data_types
```

2313 **3.9.3.11.2 Grammar**

```
data_types:  
  <tosca datatype def 1>  
  ...  
  <tosca datatype def n>
```

2314 **3.9.3.11.3 Example**

```
data_types:  
  # A complex datatype definition  
  simple_contactinfo_type:  
    properties:  
      name:  
        type: string  
      email:  
        type: string  
      phone:  
        type: string  
  
  # datatype definition derived from an existing type  
  full_contact_info:  
    derived_from: simple_contact_info  
    properties:  
      street_address:  
        type: string
```

```
city:
  type: string
state:
  type: string
postalcode:
  type: string
```

2315 **3.9.3.12 capability_types**

2316 This optional keyname lists the Capability Types that provide the reusable type definitions that can be
2317 used to describe features Node Templates or Node Types can declare they support.

2318 **3.9.3.12.1 Keyname**

```
capability_types
```

2319 **3.9.3.12.2 Grammar**

```
capability_types:
  <capability_type_defn 1>
  ...
  <capability_type_defn n>
```

2320 **3.9.3.12.3 Example**

```
capability_types:
  mycompany.mytypes.myCustomEndpoint:
    derived_from: toska.capabilities.Endpoint
    properties:
      # more details ...

  mycompany.mytypes.myCustomFeature:
    derived_from: toska.capabilities.Feature
    properties:
      # more details ...
```

2321 **3.9.3.13 interface_types**

2322 This optional keyname lists the Interface Types that provide the reusable type definitions that can be used
2323 to describe operations for on TOSCA entities such as Relationship Types and Node Types.

2324 **3.9.3.13.1 Keyname**

```
interface_types
```

2325 **3.9.3.13.2 Grammar**

```
interface_types:
  <interface_type_defn 1>
  ...
  <interface_type_defn n>
```

2326 **3.9.3.13.3 Example**

```
interface_types:
```

```
mycompany.interfaces.service.Signal:
  signal_begin_receive:
    description: Operation to signal start of some message processing.
  signal_end_receive:
    description: Operation to signal end of some message processed.
```

2327 **3.9.3.14 relationship_types**

2328 This optional keyname lists the Relationship Types that provide the reusable type definitions that can be
2329 used to describe dependent relationships between Node Templates or Node Types.

2330 **3.9.3.14.1 Keyname**

```
relationship_types
```

2331 **3.9.3.14.2 Grammar**

```
relationship_types:
  <relationship type defn 1>
  ...
  <relationship type defn n>
```

2332 **3.9.3.14.3 Example**

```
relationship_types:
  mycompany.mytypes.myCustomClientServerType:
    derived_from: toasca.relationships.HostedOn
    properties:
      # more details ...

  mycompany.mytypes.myCustomConnectionType:
    derived_from: toasca.relationships.ConnectsTo
    properties:
      # more details ...
```

2333 **3.9.3.15 node_types**

2334 This optional keyname lists the Node Types that provide the reusable type definitions for software
2335 components that Node Templates can be based upon.

2336 **3.9.3.15.1 Keyname**

```
node_types
```

2337 **3.9.3.15.2 Grammar**

```
node_types:
  <node type defn 1>
  ...
  <node type defn n>
```

2338 **3.9.3.15.3 Example**

```
node_types:
  my_webapp_node_type:
```

```
derived_from: WebApplication
properties:
  my_port:
    type: integer

my_database_node_type:
  derived_from: Database
  capabilities:
    mytypes.myfeatures.transactSQL
```

2339 **3.9.3.15.4 Notes**

- 2340 • The node types listed as part of the **node_types** block can be mapped to the list of **NodeType**
2341 definitions as described by the TOSCA v1.0 specification.

2342 **3.9.3.16 group_types**

2343 This optional keyname lists the Group Types that are defined by this Service Template.

2344 **3.9.3.16.1 Keyname**

```
group_types
```

2345 **3.9.3.16.2 Grammar**

```
group_types:
  <group type defn 1>
  ...
  <group type defn n>
```

2346 **3.9.3.16.3 Example**

```
group_types:
  mycompany.mytypes.myScalingGroup:
    derived_from: toasca.groups.Root
```

2347 **3.9.3.17 policy_types**

2348 This optional keyname lists the Policy Types that are defined by this Service Template.

2349 **3.9.3.17.1 Keyname**

```
policy_types
```

2350 **3.9.3.17.2 Grammar**

```
policy_types:
  <policy type defn 1>
  ...
  <policy type defn n>
```

2351 **3.9.3.17.3 Example**

```
policy_types:
  mycompany.mytypes.myScalingPolicy:
```

derived_from: toska.policies.Scaling

2352

4 TOSCA functions

2353

Except for the examples, this section is **normative** and includes functions that are supported for use within a TOSCA Service Template.

2354

2355

4.1 Reserved Function Keywords

2356

The following keywords MAY be used in some TOSCA function in place of a TOSCA Node or Relationship Template name. A TOSCA orchestrator will interpret them at the time the function would be evaluated at runtime as described in the table below. Note that some keywords are only valid in the context of a certain TOSCA entity as also denoted in the table.

2357

2358

2359

2360

Keyword	Valid Contexts	Description
SELF	Node Template or Relationship Template	A TOSCA orchestrator will interpret this keyword as the Node or Relationship Template instance that contains the function at the time the function is evaluated.
SOURCE	Relationship Template only.	A TOSCA orchestrator will interpret this keyword as the Node Template instance that is at the source end of the relationship that contains the referencing function.
TARGET	Relationship Template only.	A TOSCA orchestrator will interpret this keyword as the Node Template instance that is at the target end of the relationship that contains the referencing function.
HOST	Node Template only	A TOSCA orchestrator will interpret this keyword to refer to the all nodes that “host” the node using this reference (i.e., as identified by its HostedOn relationship). Specifically, TOSCA orchestrators that encounter this keyword when evaluating the get_attribute or get_property functions SHALL search each node along the “HostedOn” relationship chain starting at the immediate node that hosts the node where the function was evaluated (and then that node’s host node, and so forth) until a match is found or the “HostedOn” relationship chain ends.

2361

2362

4.2 Environment Variable Conventions

2363

4.2.1 Reserved Environment Variable Names and Usage

2364

TOSCA orchestrators utilize certain reserved keywords in the execution environments that implementation artifacts for Node or Relationship Templates operations are executed in. They are used to provide information to these implementation artifacts such as the results of TOSCA function evaluation or information about the instance model of the TOSCA application

2365

2366

2367

2368

2369

The following keywords are reserved environment variable names in any TOSCA supported execution environment:

2370

Keyword	Valid Contexts	Description
TARGETS	Relationship Template only.	<ul style="list-style-type: none"> For an implementation artifact that is executed in the context of a relationship, this keyword, if present, is used to supply a list of Node Template instances in a TOSCA application's instance model that are currently target of the context relationship. The value of this environment variable will be a comma-separated list of identifiers of the single target node instances (i.e., the tosca_id attribute of the node).
TARGET	Relationship Template only.	<ul style="list-style-type: none"> For an implementation artifact that is executed in the context of a relationship, this keyword, if present, identifies a Node Template instance in a TOSCA application's instance model that is a target of the context relationship, and which is being acted upon in the current operation. The value of this environment variable will be the identifier of the single target node instance (i.e., the tosca_id attribute of the node).
SOURCES	Relationship Template only.	<ul style="list-style-type: none"> For an implementation artifact that is executed in the context of a relationship, this keyword, if present, is used to supply a list of Node Template instances in a TOSCA application's instance model that are currently source of the context relationship. The value of this environment variable will be a comma-separated list of identifiers of the single source node instances (i.e., the tosca_id attribute of the node).
SOURCE	Relationship Template only.	<ul style="list-style-type: none"> For an implementation artifact that is executed in the context of a relationship, this keyword, if present, identifies a Node Template instance in a TOSCA application's instance model that is a source of the context relationship, and which is being acted upon in the current operation. The value of this environment variable will be the identifier of the single source node instance (i.e., the tosca_id attribute of the node).

2371

2372 For scripts (or implementation artifacts in general) that run in the context of relationship operations, select
 2373 properties and attributes of both the relationship itself as well as select properties and attributes of the
 2374 source and target node(s) of the relationship can be provided to the environment by declaring respective
 2375 operation inputs.

2376

2377 Declared inputs from mapped properties or attributes of the source or target node (selected via the
 2378 **SOURCE** or **TARGET** keyword) will be provided to the environment as variables having the exact same name
 2379 as the inputs. In addition, the same values will be provided for the complete set of source or target nodes,
 2380 however prefixed with the ID if the respective nodes. By means of the **SOURCES** or **TARGETS** variables
 2381 holding the complete set of source or target node IDs, scripts will be able to iterate over corresponding
 2382 inputs for each provided ID prefix.

2383

2384 The following example snippet shows an imaginary relationship definition from a load-balancer node to
 2385 worker nodes. A script is defined for the **add_target** operation of the Configure interface of the
 2386 relationship, and the **ip_address** attribute of the target is specified as input to the script:

2387

```
node_templates:
  load_balancer:
    type: some.vendor.LoadBalancer
    requirements:
      - member:
```



```

relationship: some.vendor.LoadBalancerToMember
interfaces:
  Configure:
    add_target:
      inputs:
        member_ip: { get_attribute: [ TARGET, ip_address ] }
      implementation: scripts/configure_members.py

```

2388 The **add_target** operation will be invoked, whenever a new target member is being added to the load-
 2389 balancer. With the above inputs declaration, a **member_ip** environment variable that will hold the IP
 2390 address of the target being added will be provided to the **configure_members.py** script. In addition, the
 2391 IP addresses of all current load-balancer members will be provided as environment variables with a
 2392 naming scheme of **<target node ID>_member_ip**. This will allow, for example, scripts that always just
 2393 write the complete list of load-balancer members into a configuration file to do so instead of updating
 2394 existing list, which might be more complicated.

2395 Assuming that the TOSCA application instance includes five load-balancer members, **node1** through
 2396 **node5**, where **node5** is the current target being added, the following environment variables (plus
 2397 potentially more variables) would be provided to the script:

```

# the ID of the current target and the IDs of all targets
TARGET=node5
TARGETS=node1,node2,node3,node4,node5

# the input for the current target and the inputs of all targets
member_ip=10.0.0.5
node1_member_ip=10.0.0.1
node2_member_ip=10.0.0.2
node3_member_ip=10.0.0.3
node4_member_ip=10.0.0.4
node5_member_ip=10.0.0.5

```

2398 With code like shown in the snippet below, scripts could then iterate of all provided **member_ip** inputs:

```

#!/usr/bin/python
import os

targets = os.environ['TARGETS'].split(',')

for t in targets:
    target_ip = os.environ.get('%s_member_ip' % t)
    # do something with target_ip ...

```

2399 4.2.2 Prefixed vs. Unprefixed TARGET names

2400 The list target node types assigned to the TARGETS key in an execution environment would have names
 2401 prefixed by unique IDs that distinguish different instances of a node in a running model. Future drafts of
 2402 this specification will show examples of how these names/IDs will be expressed.

2403 4.2.2.1 Notes

- 2404 • Target of interest is always un-prefixed. Prefix is the target opaque ID. The IDs can be used to
 2405 find the environment var. for the corresponding target. Need an example here.
- 2406 • If you have one node that contains multiple targets this would also be used (add or remove target
 2407 operations would also use this you would get set of all current targets).

2408 4.3 Intrinsic functions

2409 These functions are supported within the TOSCA template for manipulation of template data.

2410 4.3.1 concat

2411 The **concat** function is used to concatenate two or more string values within a TOSCA service template.

2412 4.3.1.1 Grammar

```
concat: [<string_value_expressions_*> ]
```

2413 4.3.1.2 Parameters

Parameter	Required	Type	Description
<string_value_expressions_*>	yes	list of string or string value expressions	A list of one or more strings (or expressions that result in a string value) which can be concatenated together into a single string.

2414 4.3.1.3 Examples

```
outputs:  
  description: Concatenate the URL for a server from other template values  
  server_url:  
    value: { concat: [ 'http://',  
                      get_attribute: [ server, public_address ],  
                      ':',  
                      get_attribute: [ server, port ] ] }
```

2415 4.3.2 token

2416 The **token** function is used within a TOSCA service template on a string to parse out (tokenize)
2417 substrings separated by one or more token characters within a larger string.

2418 4.3.2.1 Grammar

```
token: [ <string_with_tokens>, <string_of_token_chars>, <substring_index> ]
```

2419 4.3.2.2 Parameters

Parameter	Required	Type	Description
string_with_tokens	yes	string	The composite string that contains one or more substrings separated by token characters.
string_of_token_chars	yes	string	The string that contains one or more token characters that separate substrings within the composite string.
substring_index	yes	integer	The integer indicates the index of the substring to return from the composite string. Note that the first substring is denoted by using the '0' (zero) integer value.

2420 4.3.2.3 Examples

```
outputs:
```

```

webserver_port:
  description: the port provided at the end of my server's endpoint's IP address
  value: { token: [ get_attribute: [ my_server, data_endpoint, ip_address ],
                  ':',
                  1 ] }

```

2421 4.4 Property functions

2422 These functions are used within a service template to obtain property values from property definitions
 2423 declared elsewhere in the same service template. These property definitions can appear either directly in
 2424 the service template itself (e.g., in the inputs section) or on entities (e.g., node or relationship templates)
 2425 that have been modeled within the template.

2426
 2427 Note that the **get_input** and **get_property** functions may only retrieve the static values of property
 2428 definitions of a TOSCA application as defined in the TOSCA Service Template. The **get_attribute**
 2429 function should be used to retrieve values for attribute definitions (or property definitions reflected as
 2430 attribute definitions) from the runtime instance model of the TOSCA application (as realized by the
 2431 TOSCA orchestrator).

2432 4.4.1 get_input

2433 The **get_input** function is used to retrieve the values of properties declared within the **inputs** section of
 2434 a TOSCA Service Template.

2435 4.4.1.1 Grammar

```
get_input: <input_property_name>
```

2436 4.4.1.2 Parameters

Parameter	Required	Type	Description
<input_property_name>	yes	string	The name of the property as defined in the inputs section of the service template.

2437 4.4.1.3 Examples

```

inputs:
  cpus:
    type: integer

node_templates:
  my_server:
    type: toska.nodes.Compute
    capabilities:
      host:
        properties:
          num_cpus: { get_input: cpus }

```

2438 4.4.2 get_property

2439 The **get_property** function is used to retrieve property values between modelable entities defined in the
 2440 same service template.

2441 **4.4.2.1 Grammar**

```
get_property: [ <modelable_entity_name>, <optional_req_or_cap_name>,
<property_name>, <nested_property_name_or_index_1>, ...,
<nested_property_name_or_index_n> ]
```

2442 **4.4.2.2 Parameters**

Parameter	Required	Type	Description
<modelable_entity_name> SELF SOURCE TARGET HOST	yes	string	The required name of a modelable entity (e.g., Node Template or Relationship Template name) as declared in the service template that contains the named property definition the function will return the value from. See section B.1 for valid keywords.
<optional_req_or_cap_name>	no	string	The optional name of the requirement or capability name within the modelable entity (i.e., the <modelable_entity_name> which contains the named property definition the function will return the value from. Note: If the property definition is located in the modelable entity directly, then this parameter MAY be omitted.
<property_name>	yes	string	The name of the property definition the function will return the value from.
<nested_property_name_or_index_*>	no	string integer	Some TOSCA properties are complex (i.e., composed as nested structures). These parameters are used to dereference into the names of these nested structures when needed. Some properties represent list types. In these cases, an index may be provided to reference a specific entry in the list (as named in the previous parameter) to return.

2443 **4.4.2.3 Examples**

2444 The following example shows how to use the **get_property** function with an actual Node Template
2445 name:

```
node_templates:
  mysql_database:
    type: toasca.nodes.Database
    properties:
      name: sql_database1
  wordpress:
    type: toasca.nodes.WebApplication.WordPress
    ...
    interfaces:
      Standard:
        configure:
          inputs:
            wp_db_name: { get_property: [ mysql_database, name ] }
```

2446 The following example shows how to use the **get_property** function using the SELF keyword:

```
node_templates:
```

```

mysql_database:
  type: toska.nodes.Database
  ...
  capabilities:
    database_endpoint:
      properties:
        port: 3306

wordpress:
  type: toska.nodes.WebApplication.WordPress
  requirements:
    ...
    - database_endpoint: mysql_database
  interfaces:
    Standard:
      create: wordpress_install.sh
      configure:
        implementation: wordpress_configure.sh
      inputs:
        ...
        wp_db_port: { get_property: [ SELF, database_endpoint, port ] }

```

2447 The following example shows how to use the `get_property` function using the `TARGET` keyword:

```

relationship_templates:
  my_connection:
    type: ConnectsTo
    interfaces:
      Configure:
        inputs:
          targets_value: { get_property: [ TARGET, value ] }

```

2448 4.5 Attribute functions

2449 These functions (attribute functions) are used within an instance model to obtain attribute values from
 2450 instances of nodes and relationships that have been created from an application model described in a
 2451 service template. The instances of nodes or relationships can be referenced by their name as assigned
 2452 in the service template or relative to the context where they are being invoked.

2453 4.5.1 `get_attribute`

2454 The `get_attribute` function is used to retrieve the values of named attributes declared by the
 2455 referenced node or relationship template name.

2456 4.5.1.1 Grammar

```

get_attribute: [ <modelable_entity_name>, <optional_req_or_cap_name>,
<attribute_name>, <nested_attribute_name_or_index_1>, ...,
<nested_attribute_name_or_index_n> ]

```

2457 **4.5.1.2 Parameters**

Parameter	Required	Type	Description
<modelable_entity_name> SELF SOURCE TARGET HOST	yes	string	The required name of a modelable entity (e.g., Node Template or Relationship Template name) as declared in the service template that contains the named attribute definition the function will return the value from. See section B.1 for valid keywords.
<optional_req_or_cap_name>	no	string	The optional name of the requirement or capability name within the modelable entity (i.e., the <modelable_entity_name> which contains the named attribute definition the function will return the value from. Note: If the attribute definition is located in the modelable entity directly, then this parameter MAY be omitted.
<attribute_name>	yes	string	The name of the attribute definition the function will return the value from.
<nested_attribute_name_or_index_*>	no	string integer	Some TOSCA attributes are complex (i.e., composed as nested structures). These parameters are used to dereference into the names of these nested structures when needed. Some attributes represent list types. In these cases, an index may be provided to reference a specific entry in the list (as named in the previous parameter) to return.

2458 **4.5.1.3 Examples:**

2459 The attribute functions are used in the same way as the equivalent Property functions described above.
2460 Please see their examples and replace “get_property” with “get_attribute” function name.

2461 **4.5.1.4 Notes**

2462 These functions are used to obtain attributes from instances of node or relationship templates by the
2463 names they were given within the service template that described the application model (pattern).

- 2464 • These functions only work when the orchestrator can resolve to a single node or relationship
2465 instance for the named node or relationship. This essentially means this is acknowledged to work
2466 only when the node or relationship template being referenced from the service template has a
2467 cardinality of 1 (i.e., there can only be one instance of it running).

2468 **4.6 Operation functions**

2469 These functions are used within an instance model to obtain values from interface operations. These can
2470 be used in order to set an attribute of a node instance at runtime or to pass values from one operation to
2471 another.

2472 **4.6.1 get_operation_output**

2473 The **get_operation_output** function is used to retrieve the values of variables exposed / exported from
2474 an interface operation.

2475 **4.6.1.1 Grammar**

```
get_operation_output: <modelable_entity_name>, <interface_name>, <operation_name>,
<output_variable_name>
```

2476 **4.6.1.2 Parameters**

Parameter	Required	Type	Description
<modelable_entity_name> SELF SOURCE TARGET	yes	string	The required name of a modelable entity (e.g., Node Template or Relationship Template name) as declared in the service template that implements the named interface and operation.
<interface_name>	Yes	string	The required name of the interface which defines the operation.
<operation_name>	yes	string	The required name of the operation whose value we would like to retrieve.
<output_variable_name>	Yes	string	The required name of the variable that is exposed / exported by the operation.

2477 **4.6.1.3 Notes**

- If operation failed, then ignore its outputs. Orchestrators should allow orchestrators to continue running when possible past deployment in the lifecycle. For example, if an update fails, the application should be allowed to continue running and some other method would be used to alert administrators of the failure.

2482 **4.7 Navigation functions**

- This version of the TOSCA Simple Profile does not define any model navigation functions.

2484 **4.7.1 get_nodes_of_type**

2485 The `get_nodes_of_type` function can be used to retrieve a list of all known instances of nodes of the
2486 declared Node Type.

2487 **4.7.1.1 Grammar**

```
get_nodes_of_type: <node_type_name>
```

2488 **4.7.1.2 Parameters**

Parameter	Required	Type	Description
<node_type_name>	yes	string	The required name of a Node Type that a TOSCA orchestrator would use to search a running application instance in order to return all unique, named node instances of that type.

2489 **4.7.1.3 Returns**

Return Key	Type	Description
TARGETS	<see above>	The list of node instances from the current application instance that match the node_type_name supplied as an input parameter of this function.

2490 **4.8 Artifact functions**

2491 **4.8.1 get_artifact**

2492 The `get_artifact` function is used to retrieve artifact location between modelable entities defined in the
2493 same service template.

2494 **4.8.1.1 Grammar**

```
get_artifact: [ <modelable_entity_name>, <artifact_name>, <location>, <remove> ]
```

2495 **4.8.1.2 Parameters**

Parameter	Required	Type	Description
<modelable entity name> SELF SOURCE TARGET HOST	yes	string	The required name of a modelable entity (e.g., Node Template or Relationship Template name) as declared in the service template that contains the named property definition the function will return the value from. See section B.1 for valid keywords.
<artifact_name>	yes	string	The name of the artifact definition the function will return the value from.
<location> LOCAL_FILE	no	string	Location value must be either a valid path e.g. '/etc/var/my_file' or 'LOCAL_FILE'. If the value is LOCAL_FILE the orchestrator is responsible for providing a path as the result of the <code>get_artifact</code> call where the artifact file can be accessed. The orchestrator will also remove the artifact from this location at the end of the operation. If the location is a path specified by the user the orchestrator is responsible to copy the artifact to the specified location. The orchestrator will return the path as the value of the <code>get_artifact</code> function and leave the file here after the execution of the operation.
remove	no	boolean	Boolean flag to override the orchestrator default behavior so it will remove or not the artifact at the end of the operation execution. If not specified the removal will depends of the location e.g. removes it in case of 'LOCAL_FILE' and keeps it in case of a path. If true the artifact will be removed by the orchestrator at the end of the operation execution, if false it will not be removed.

2496 **4.8.1.3 Examples**

2497 The following example uses a snippet of a WordPress [WordPress] web application to show how to use
2498 the `get_artifact` function with an actual Node Template name:

2499 **4.8.1.3.1 Example: Retrieving artifact without specified location**

```
node_templates:
  wordpress:
    type: tosca.nodes.WebApplication.WordPress
    ...
  interfaces:
```



```

Standard:
  configure:
    create:
      implementation: wordpress_install.sh
      inputs
        wp_zip: { get_artifact: [ SELF, zip ] }
artifacts:
  zip: /data/wordpress.zip

```

2500 In such implementation the TOSCA orchestrator may provide the **wordpress.zip** archive as

- 2501 • a local URL (example: <file:///home/user/wordpress.zip>) or
- 2502 • a remote one (example: <http://cloudrepo:80/files/wordpress.zip>) where some orchestrator
- 2503 may indeed provide some global artifact repository management features.

2504 4.8.1.3.2 Example: Retrieving artifact as a local path

2505 The following example explains how to force the orchestrator to copy the file locally before calling the
 2506 operation's implementation script:

2507

```

node_templates:

wordpress:
  type: tosca.nodes.WebApplication.WordPress
  ...
  interfaces:
    Standard:
      configure:
        create:
          implementation: wordpress_install.sh
          inputs
            wp_zip: { get_artifact: [ SELF, zip, LOCAL_FILE ] }
  artifacts:
    zip: /data/wordpress.zip

```

2508 In such implementation the TOSCA orchestrator must provide the **wordpress.zip** archive as a local path
 2509 (example: </tmp/wordpress.zip>) and **will remove it** after the operation is completed.

2510 4.8.1.3.3 Example: Retrieving artifact in a specified location

2511 The following example explains how to force the orchestrator to copy the file locally to a specific location
 2512 before calling the operation's implementation script :

2513

```

node_templates:

wordpress:
  type: tosca.nodes.WebApplication.WordPress
  ...
  interfaces:
    Standard:
      configure:
        create:
          implementation: wordpress_install.sh
          inputs
            wp_zip: { get_artifact: [ SELF, zip, C:/wpdata/wp.zip ] }

```

```
artifacts:  
  zip: /data/wordpress.zip
```

2514 In such implementation the TOSCA orchestrator must provide the wordpress.zip archive as a local path
2515 (example: C:/wpdata/wp.zip) and **will let it** after the operation is completed.

2516 **4.9 Context-based Entity names (global)**

2517 Future versions of this specification will address methods to access entity names based upon the context
2518 in which they are declared or defined.

2519 **4.9.1.1 Goals**

- 2520 • Using the full paths of modelable entity names to qualify context with the future goal of a more
2521 robust get_attribute function: e.g., get_attribute(<context-based-entity-name>, <attribute name>)

2522

5 TOSCA normative type definitions

2523

Except for the examples, this section is **normative** and contains normative type definitions which must be supported for conformance to this specification.

2524

2525

2526

The declarative approach is heavily dependent of the definition of basic types that a declarative container must understand. The definition of these types must be very clear such that the operational semantics can be precisely followed by a declarative container to achieve the effects intended by the modeler of a topology in an interoperable manner.

2527

2528

2529

2530

5.1 Assumptions

2531

- Assumes alignment with/dependence on XML normative types proposal for TOSCA v1.1

2532

- Assumes that the normative types will be versioned and the TOSCA TC will preserve backwards compatibility.

2533

2534

- Assumes that security and access control will be addressed in future revisions or versions of this specification.

2535

2536

5.2 TOSCA normative type names

2537

Every normative type has three names declared:

2538

1. **Type URI** – This is the unique identifying name for the type.

2539

- a. These are reserved names within the TOSCA namespace.

2540

2. **Shorthand Name** – This is the shorter (simpler) name that can be used in place of its corresponding, full **Type URI** name.

2541

2542

- a. These are reserved names within TOSCA namespace that MAY be used in place of the full Type URI.

2543

2544

- b. Profiles of the OASIS TOSCA Simple Profile specification SHALL assure non-collision of names for new types when they are introduced.

2545

2546

- c. TOSCA type designers SHOULD NOT create new types with names that would collide with any TOSCA normative type Shorthand Name.

2547

2548

3. **Type Qualified Name** – This is a modified **Shorthand Name** that includes the “**tosca:**” namespace prefix which clearly qualifies it as being part of the TOSCA namespace.

2549

2550

- a. This name MAY be used to assure there is no collision when types are imported from other (non) TOSCA approved sources.

2551

2552

5.2.1 Additional requirements

2553

- **Case sensitivity** - TOSCA Type URI, Shorthand and Type Qualified names SHALL be treated as case sensitive.

2554

2555

- The case of each type name has been carefully selected by the TOSCA working group and TOSCA orchestrators and processors SHALL strictly recognize the name casing as specified in this specification or any of its approved profiles.

2556

2557

2558 **5.3 Data Types**

2559 **5.3.1 tosca.datatypes.Root**

2560 This is the default (root) TOSCA Root Type definition that all complex TOSCA Data Types derive from.

2561 **5.3.1.1 Definition**

2562 The TOSCA Root type is defined as follows:

```
tosca.datatypes.Root:
  description: The TOSCA root Data Type all other TOSCA base Data Types derive from
```

2563 **5.3.2 tosca.datatypes.Credential**

2564 The Credential type is a complex TOSCA data Type used when describing authorization credentials used
2565 to access network accessible resources.

Shorthand Name	Credential
Type Qualified Name	tosca:Credential
Type URI	tosca.datatypes.Credential

2566 **5.3.2.1 Properties**

Name	Required	Type	Constraints	Description
protocol	no	string	None	The optional protocol name.
token_type	yes	string	default: password	The required token type.
token	yes	string	None	The required token used as a credential for authorization or access to a networked resource.
keys	no	map of string	None	The optional list of protocol-specific keys or assertions.
user	no	string	None	The optional user (name or ID) used for non-token based credentials.

2567 **5.3.2.2 Definition**

2568 The TOSCA Credential type is defined as follows:

```
tosca.datatypes.Credential:
  derived_from: tosca.datatypes.Root
  properties:
    protocol:
      type: string
      required: false
    token_type:
      type: string
      default: password
    token:
      type: string
    keys:
      type: map
      required: false
```

```
entry_schema:
  type: string
user:
  type: string
  required: false
```

2569 5.3.2.3 Additional requirements

- 2570
- TOSCA Orchestrators SHALL interpret and validate the value of the **token** property based upon the value of the **token_type** property.
- 2571

2572 5.3.2.4 Notes

- 2573
- Specific token types and encoding them using network protocols are not defined or covered in this specification.
- 2574
- The use of transparent user names (IDs) or passwords are not considered best practice.
- 2575

2576 5.3.2.5 Examples

2577 5.3.2.5.1 Provide a simple user name and password without a protocol or 2578 standardized token format

```
<some_tosca_entity>:
  properties:
    my_credential:
      type: Credential
      properties:
        user: myusername
        token: mypassword
```

2579 5.3.2.5.2 HTTP Basic access authentication credential

```
<some_tosca_entity>:
  properties:
    my_credential: # type: Credential
      protocol: http
      token_type: basic_auth
      # Username and password are combined into a string
      # Note: this would be base64 encoded before transmission by any impl.
      token: myusername:mypassword
```

2580 5.3.2.5.3 X-Auth-Token credential

```
<some_tosca_entity>:
  properties:
    my_credential: # type: Credential
      protocol: xauth
      token_type: X-Auth-Token
      # token encoded in Base64
      token: 604bbe45ac7143a79e14f3158df67091
```

2581 5.3.2.5.4 OAuth bearer token credential

```
<some_tosca_entity>:
```

```

properties:
  my_credential: # type: Credential
    protocol: oauth2
    token_type: bearer
    # token encoded in Base64
    token: 8ao9nE2DEjr1zCsicWmpBC

```

2582 **5.3.2.6 OpenStack SSH Keypair**

```

<some_tosca_entity>:
  properties:
    my_ssh_keypair: # type: Credential
      protocol: ssh
      token_type: identifier
      # token is a reference (ID) to an existing keypair (already installed)
      token: <keypair_id>

```

2583

2584 **5.3.3 tosca.datatypes.TimeInterval**

2585 The TimeInterval type is a complex TOSCA data Type used when describing a period of time using the
 2586 YAML ISO 8601 format to declare the start and end times.

Shorthand Name	TimeInterval
Type Qualified Name	tosca:TimeInterval
Type URI	tosca.datatypes.TimeInterval

2587 **5.3.3.1 Properties**

Name	Required	Type	Constraints	Description
start_time	yes	timestamp	None	The inclusive start time for the time interval.
end_time	yes	timestamp	None	The inclusive end time for the time interval.

2588 **5.3.3.2 Definition**

2589 The TOSCA TimeInterval type is defined as follows:

```

tosca.datatypes.TimeInterval:
  derived_from: tosca.datatypes.Root
  properties:
    start_time:
      type: timestamp
      required: true
    end_time:
      type: timestamp
      required: true

```

2590 **5.3.3.3 Examples**

2591 **5.3.3.3.1 Multi-day evaluation time period**

```

properties:
  description:

```

```

evaluation_period: Evaluate a service for a 5-day period across time zones
type: TimeInterval
start_time: 2016-04-04-15T00:00:00Z
end_time: 2016-04-08T21:59:43.10-06:00

```

2592 5.3.4 **tosca.datatypes.network.NetworkInfo**

2593 The Network type is a complex TOSCA data type used to describe logical network information.

Shorthand Name	NetworkInfo
Type Qualified Name	tosca:NetworkInfo
Type URI	tosca.datatypes.network.NetworkInfo

2594 5.3.4.1 Properties

Name	Type	Constraints	Description
network_name	string	None	The name of the logical network. e.g., "public", "private", "admin". etc.
network_id	string	None	The unique ID of for the network generated by the network provider.
addresses	string []	None	The list of IP addresses assigned from the underlying network.

2595 5.3.4.2 Definition

2596 The TOSCA NetworkInfo data type is defined as follows:

```

tosca.datatypes.network.NetworkInfo:
  derived_from: toska.datatypes.Root
  properties:
    network_name:
      type: string
    network_id:
      type: string
    addresses:
      type: list
      entry_schema:
        type: string

```

2597 5.3.4.3 Examples

2598 Example usage of the NetworkInfo data type:

```

<some_tosca_entity>:
  properties:
    private_network:
      network_name: private
      network_id: 3e54214f-5c09-1bc9-9999-44100326da1b
      addresses: [ 10.111.128.10 ]

```

2599 5.3.4.4 Additional Requirements

- 2600 • It is expected that TOSCA orchestrators **MUST** be able to map the **network_name** from the
- 2601 TOSCA model to underlying network model of the provider.

- The properties (or attributes) of NetworkInfo may or may not be required depending on usage context.

2604 5.3.5 tosca.datatypes.network.PortInfo

2605 The PortInfo type is a complex TOSCA data type used to describe network port information.

Shorthand Name	PortInfo
Type Qualified Name	tosca:PortInfo
Type URI	tosca.datatypes.network.PortInfo

2606 5.3.5.1 Properties

Name	Type	Constraints	Description
port_name	string	None	The logical network port name.
port_id	string	None	The unique ID for the network port generated by the network provider.
network_id	string	None	The unique ID for the network.
mac_address	string	None	The unique media access control address (MAC address) assigned to the port.
addresses	string []	None	The list of IP address(es) assigned to the port.

2607 5.3.5.2 Definition

2608 The TOSCA PortInfo type is defined as follows:

```

tosca.datatypes.network.PortInfo:
  derived_from: tosca.datatypes.Root
  properties:
    port_name:
      type: string
    port_id:
      type: string
    network_id:
      type: string
    mac_address:
      type: string
    addresses:
      type: list
      entry_schema:
        type: string

```

2609 5.3.5.3 Examples

2610 Example usage of the PortInfo data type:

```

ethernet_port:
  properties:
    port_name: port1
    port_id: 2c0c7a37-691a-23a6-7709-2d10ad041467
    network_id: 3e54214f-5c09-1bc9-9999-44100326da1b
    mac_address: f1:18:3b:41:92:1e
    addresses: [ 172.24.9.102 ]

```


2611 **5.3.5.4 Additional Requirements**

- 2612 • It is expected that TOSCA orchestrators MUST be able to map the `port_name` from the TOSCA
- 2613 model to underlying network model of the provider.
- 2614 • The properties (or attributes) of PortInfo may or may not be required depending on usage context.

2615 **5.3.6 tosca.datatypes.network.PortDef**

2616 The PortDef type is a TOSCA data Type used to define a network port.

Shorthand Name	PortDef
Type Qualified Name	tosca:PortDef
Type URI	tosca.datatypes.network.PortDef

2617 **5.3.6.1 Definition**

2618 The TOSCA PortDef type is defined as follows:

```
tosca.datatypes.network.PortDef:  
  derived_from: integer  
  constraints:  
    - in_range: [ 1, 65535 ]
```

2619 **5.3.6.2 Examples**

2620 Simple usage of a PortDef property type:

```
properties:  
  listen_port: 9090
```

2621 Example declaration of a property for a custom type based upon PortDef:

```
properties:  
  listen_port:  
    type: PortDef  
    default: 9000  
    constraints:  
      - in_range: [ 9000, 9090 ]
```

2622 **5.3.7 tosca.datatypes.network.PortSpec**

2623 The PortSpec type is a complex TOSCA data Type used when describing port specifications for a
2624 network connection.

Shorthand Name	PortSpec
Type Qualified Name	tosca:PortSpec
Type URI	tosca.datatypes.network.PortSpec

2625 **5.3.7.1 Properties**

Name	Required	Type	Constraints	Description
protocol	yes	string	default: tcp	The required protocol used on the port.
source	no	PortDef	See PortDef	The optional source port.

Name	Required	Type	Constraints	Description
source_range	no	range	in_range: [1, 65536]	The optional range for source port.
target	no	PortDef	See PortDef	The optional target port.
target_range	no	range	in_range: [1, 65536]	The optional range for target port.

2626 5.3.7.2 Definition

2627 The TOSCA PortSpec type is defined as follows:

```
tosca.datatypes.network.PortSpec:
  derived_from: tosca.datatypes.Root
  properties:
    protocol:
      type: string
      required: true
      default: tcp
      constraints:
        - valid_values: [ udp, tcp, igmp ]
    target:
      type: PortDef
      required: false
    target_range:
      type: range
      required: false
      constraints:
        - in_range: [ 1, 65535 ]
    source:
      type: PortDef
      required: false
    source_range:
      type: range
      required: false
      constraints:
        - in_range: [ 1, 65535 ]
```

2628 5.3.7.3 Additional requirements

- 2629
- 2630
- 2631
- 2632
- 2633
- 2634
- A valid PortSpec MUST have at least one of the following properties: **target**, **target_range**, **source** or **source_range**.
 - A valid PortSpec MUST have a value for the **source** property that is within the numeric range specified by the property **source_range** when **source_range** is specified.
 - A valid PortSpec MUST have a value for the **target** property that is within the numeric range specified by the property **target_range** when **target_range** is specified.

2635 5.3.7.4 Examples

2636 Example usage of the PortSpec data type:

```
# example properties in a node template
some_endpoint:
  properties:
    ports:
      user_port:
        protocol: tcp
```

```
target: 50000
target_range: [ 20000, 60000 ]
source: 9000
source_range: [ 1000, 10000 ]
```

2637 5.4 Artifact Types

2638 TOSCA Artifacts Types represent the types of packages and files used by the orchestrator when
2639 deploying TOSCA Node or Relationship Types or invoking their interfaces. Currently, artifacts are
2640 logically divided into three categories:

- 2641
- 2642 • **Deployment Types:** includes those artifacts that are used during deployment (e.g., referenced
2643 on create and install operations) and include packaging files such as RPMs, ZIPs, or TAR files.
 - 2644 • **Implementation Types:** includes those artifacts that represent imperative logic and are used to
2645 implement TOSCA Interface operations. These typically include scripting languages such as
2646 Bash (.sh), Chef [[Chef](#)] and Puppet [[Puppet](#)].
 - 2647 • **Runtime Types:** includes those artifacts that are used during runtime by a service or component
2648 of the application. This could include a library or language runtime that is needed by an
2649 application such as a PHP or Java library.

2650

2651 **Note:** Additional TOSCA Artifact Types will be developed in future drafts of this specification.

2652 5.4.1 `tosca.artifacts.Root`

2653 This is the default (root) TOSCA [Artifact Type](#) definition that all other TOSCA base Artifact Types derive
2654 from.

2655 5.4.1.1 Definition

```
tosca.artifacts.Root:
  description: The TOSCA Artifact Type all other TOSCA Artifact Types derive from
```

2656 5.4.2 `tosca.artifacts.File`

2657 This artifact type is used when an artifact definition needs to have its associated file simply treated as a
2658 file and no special handling/handlers are invoked (i.e., it is not treated as either an implementation or
2659 deployment artifact type).

Shorthand Name	File
Type Qualified Name	tosca:File
Type URI	tosca.artifacts.File

2660 5.4.2.1 Definition

```
tosca.artifacts.File:
  derived_from: tosca.artifacts.Root
```

2661 **5.4.3 Deployment Types**

2662 **5.4.3.1 `tosca.artifacts.Deployment`**

2663 This artifact type represents the parent type for all deployment artifacts in TOSCA. This class of artifacts
2664 typically represents a binary packaging of an application or service that is used to install/create or deploy
2665 it as part of a node's lifecycle.

2666 **5.4.3.1.1 Definition**

```
tosca.artifacts.Deployment:
  derived_from: tosca.artifacts.Root
  description: TOSCA base type for deployment artifacts
```

2667 **5.4.3.2 Additional Requirements**

- 2668
 - TOSCA Orchestrators MAY throw an error if it encounters a non-normative deployment artifact
2669 type that it is not able to process.

2670 **5.4.3.3 `tosca.artifacts.Deployment.Image`**

2671 This artifact type represents a parent type for any "image" which is an opaque packaging of a TOSCA
2672 Node's deployment (whether real or virtual) whose contents are typically already installed and pre-
2673 configured (i.e., "stateful") and prepared to be run on a known target container.

Shorthand Name	Deployment.Image
Type Qualified Name	tosca:Deployment.Image
Type URI	tosca.artifacts.Deployment.Image

2674 **5.4.3.3.1 Definition**

```
tosca.artifacts.Deployment.Image:
  derived_from: tosca.artifacts.Deployment
```

2675 **5.4.3.4 `tosca.artifacts.Deployment.Image.VM`**

2676 This artifact represents the parent type for all Virtual Machine (VM) image and container formatted
2677 deployment artifacts. These images contain a stateful capture of a machine (e.g., server) including
2678 operating system and installed software along with any configurations and can be run on another
2679 machine using a hypervisor which virtualizes typical server (i.e., hardware) resources.

2680 **5.4.3.4.1 Definition**

```
tosca.artifacts.Deployment.Image.VM:
  derived_from: tosca.artifacts.Deployment.Image
  description: Virtual Machine (VM) Image
```

2681 **5.4.3.4.2 Notes**

- 2682
 - Future drafts of this specification may include popular standard VM disk image (e.g., ISO, VMI,
2683 VMDX, QCOW2, etc.) and container (e.g., OVF, bare, etc.) formats. These would include
2684 consideration of disk formats such as:

2685 **5.4.4 Implementation Types**

2686 **5.4.4.1 tosca.artifacts.Implementation**

2687 This artifact type represents the parent type for all implementation artifacts in TOSCA. These artifacts are
2688 used to implement operations of TOSCA interfaces either directly (e.g., scripts) or indirectly (e.g., config.
2689 files).

2690 **5.4.4.1.1 Definition**

```
tosca.artifacts.Implementation:
  derived_from: tosca.artifacts.Root
  description: TOSCA base type for implementation artifacts
```

2691 **5.4.4.2 Additional Requirements**

- 2692
 - TOSCA Orchestrators **MAY** throw an error if it encounters a non-normative implementation
2693 artifact type that it is not able to process.

2694 **5.4.4.3 tosca.artifacts.Implementation.Bash**

2695 This artifact type represents a Bash script type that contains Bash commands that can be executed on
2696 the Unix Bash shell.

Shorthand Name	Bash
Type Qualified Name	tosca:Bash
Type URI	tosca.artifacts.Implementation.Bash

2697 **5.4.4.3.1 Definition**

```
tosca.artifacts.Implementation.Bash:
  derived_from: tosca.artifacts.Implementation
  description: Script artifact for the Unix Bash shell
  mime_type: application/x-sh
  file_ext: [ sh ]
```

2698 **5.4.4.4 tosca.artifacts.Implementation.Python**

2699 This artifact type represents a Python file that contains Python language constructs that can be executed
2700 within a Python interpreter.

Shorthand Name	Python
Type Qualified Name	tosca:Python
Type URI	tosca.artifacts.Implementation.Python

2701 **5.4.4.4.1 Definition**

```
tosca.artifacts.Implementation.Python:
  derived_from: tosca.artifacts.Implementation
  description: Artifact for the interpreted Python language
  mime_type: application/x-python
  file_ext: [ py ]
```

2702 **5.5 Capabilities Types**

2703 **5.5.1 tosca.capabilities.Root**

2704 This is the default (root) TOSCA Capability Type definition that all other TOSCA Capability Types derive
2705 from.

2706 **5.5.1.1 Definition**

```
tosca.capabilities.Root:
  description: The TOSCA root Capability Type all other TOSCA base Capability Types
  derive from
```

2707 **5.5.2 tosca.capabilities.Node**

2708 The Node capability indicates the base capabilities of a TOSCA Node Type.

Shorthand Name	Node
Type Qualified Name	tosca:Node
Type URI	tosca.capabilities.Node

2709 **5.5.2.1 Definition**

```
tosca.capabilities.Node:
  derived_from: tosca.capabilities.Root
```

2710 **5.5.3 tosca.capabilities.Compute**

2711 The Compute capability, when included on a Node Type or Template definition, indicates that the node
2712 can provide hosting on a named compute resource.

Shorthand Name	Compute
Type Qualified Name	tosca:Compute
Type URI	tosca.capabilities.Compute

2713 **5.5.3.1 Properties**

Name	Required	Type	Constraints	Description
name	no	string	None	The optional name (or identifier) of a specific compute resource for hosting.
num_cpus	no	integer	greater_or_equal: 1	Number of (actual or virtual) CPUs associated with the Compute node.
cpu_frequency	no	scalar-unit.frequency	greater_or_equal: 0.1 GHz	Specifies the operating frequency of CPU's core. This property expresses the expected frequency of one (1) CPU as provided by the property "num_cpus".
disk_size	no	scalar-unit.size	greater_or_equal: 0 MB	Size of the local disk available to applications running on the Compute node (default unit is MB).
mem_size	no	scalar-unit.size	greater_or_equal: 0 MB	Size of memory available to applications running on the Compute node (default unit is MB).

2714 **5.5.3.2 Definition**

```

tosca.capabilities.Compute:
  derived_from: tosca.capabilities.Root
  properties:
    name:
      type: string
      required: false
    num_cpus:
      type: integer
      required: false
      constraints:
        - greater_or_equal: 1
    cpu_frequency:
      type: scalar-unit.frequency
      required: false
      constraints:
        - greater_or_equal: 0.1 GHz
    disk_size:
      type: scalar-unit.size
      required: false
      constraints:
        - greater_or_equal: 0 MB
    mem_size:
      type: scalar-unit.size
      required: false
      constraints:
        - greater_or_equal: 0 MB

```

2715 **5.5.4 tosca.capabilities.Network**

2716 The Storage capability, when included on a Node Type or Template definition, indicates that the node can
 2717 provide addressibility for the resource a named network with the specified ports.

Shorthand Name	Network
Type Qualified Name	tosca:Network
Type URI	tosca.capabilities.Network

2718 **5.5.4.1 Properties**

Name	Required	Type	Constraints	Description
name	no	string	None	The otional name (or identifier) of a specific network resource.

2719 **5.5.4.2 Definition**

```

tosca.capabilities.Network:
  derived_from: tosca.capabilities.Root
  properties:
    name:
      type: string
      required: false

```

2720 **5.5.5 tosca.capabilities.Storage**

2721 The Storage capability, when included on a Node Type or Template definition, indicates that the node can
2722 provide a named storage location with specified size range.

Shorthand Name	Storage
Type Qualified Name	tosca:Storage
Type URI	tosca.capabilities.Storage

2723 **5.5.5.1 Properties**

Name	Required	Type	Constraints	Description
name	no	string	None	The optional name (or identifier) of a specific storage resource.

2724 **5.5.5.2 Definition**

```
tosca.capabilities.Storage:
  derived_from: tosca.capabilities.Root
  properties:
    name:
      type: string
      required: false
```

2725 **5.5.6 tosca.capabilities.Container**

2726 The Container capability, when included on a Node Type or Template definition, indicates that the node
2727 can act as a container for (or a host for) one or more other declared Node Types.

Shorthand Name	Container
Type Qualified Name	tosca:Container
Type URI	tosca.capabilities.Container

2728 **5.5.6.1 Properties**

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

2729 **5.5.6.2 Definition**

```
tosca.capabilities.Container:
  derived_from: tosca.capabilities.Compute
```


2730 **5.5.7 tosca.capabilities.Endpoint**

2731 This is the default TOSCA type that should be used or extended to define a network endpoint capability.
 2732 This includes the information to express a basic endpoint with a single port or a complex endpoint with
 2733 multiple ports. By default the Endpoint is assumed to represent an address on a private network unless
 2734 otherwise specified.

Shorthand Name	Endpoint
Type Qualified Name	tosca:Endpoint
Type URI	tosca.capabilities.Endpoint

2735 **5.5.7.1 Properties**

Name	Required	Type	Constraints	Description
protocol	yes	string	default: tcp	The name of the protocol (i.e., the protocol prefix) that the endpoint accepts (any OSI Layer 4-7 protocols) Examples: http, https, ftp, tcp, udp, etc.
port	no	PortDef	greater_or_equal: 1 less_or_equal: 65535	The optional port of the endpoint.
secure	no	boolean	default: false	Requests for the endpoint to be secure and use credentials supplied on the ConnectsTo relationship.
url_path	no	string	None	The optional URL path of the endpoint's address if applicable for the protocol.
port_name	no	string	None	The optional name (or ID) of the network port this endpoint should be bound to.
network_name	no	string	default: PRIVATE	The optional name (or ID) of the network this endpoint should be bound to. network_name: PRIVATE PUBLIC <network_name> <network_id>
initiator	no	string	one of: • source • target • peer default: source	The optional indicator of the direction of the connection.
ports	no	map of PortSpec	None	The optional map of ports the Endpoint supports (if more than one)

2736 **5.5.7.2 Attributes**

Name	Required	Type	Constraints	Description
ip_address	yes	string	None	Note: This is the IP address as propagated up by the associated node's host (Compute) container.

2737 **5.5.7.3 Definition**

```
tosca.capabilities.Endpoint:
  derived_from: tosca.capabilities.Root
```

```

properties:
  protocol:
    type: string
    required: true
    default: tcp
  port:
    type: PortDef
    required: false
  secure:
    type: boolean
    required: false
    default: false
  url_path:
    type: string
    required: false
  port_name:
    type: string
    required: false
  network_name:
    type: string
    required: false
    default: PRIVATE
  initiator:
    type: string
    required: false
    default: source
    constraints:
      - valid_values: [ source, target, peer ]
  ports:
    type: map
    required: false
    constraints:
      - min_length: 1
    entry_schema:
      type: PortSpec
  attributes:
    ip_address:
      type: string

```

2738 5.5.7.4 Additional requirements

- 2739
- 2740
- Although both the port and ports properties are not required, one of port or ports must be provided in a valid [Endpoint](#).

2741 5.5.8 tosca.capabilities.Endpoint.Public

2742 This capability represents a public endpoint which is accessible to the general internet (and its public IP
2743 address ranges).

2744 This public endpoint capability also can be used to create a floating (IP) address that the underlying
2745 network assigns from a pool allocated from the application's underlying public network. This floating
2746 address is managed by the underlying network such that can be routed an application's private address
2747 and remains reliable to internet clients.

Shorthand Name	Endpoint.Public
Type Qualified Name	tosca:Endpoint.Public
Type URI	tosca.capabilities.Endpoint.Public

2748 **5.5.8.1 Definition**

```
tosca.capabilities.Endpoint.Public:
  derived_from: tosca.capabilities.Endpoint
  properties:
    # Change the default network_name to use the first public network found
    network_name:
      type: string
      default: PUBLIC
      constraints:
        - equal: PUBLIC
    floating:
      description: >
        indicates that the public address should be allocated from a pool of
        floating IPs that are associated with the network.
      type: boolean
      default: false
      status: experimental
    dns_name:
      description: The optional name to register with DNS
      type: string
      required: false
      status: experimental
```

2749 **5.5.8.2 Additional requirements**

- 2750 • If the **network_name** is set to the reserved value **PRIVATE** or if the value is set to the name of
2751 network (or subnetwork) that is not public (i.e., has non-public IP address ranges assigned to it)
2752 then TOSCA Orchestrators **SHALL** treat this as an error.
- 2753 • If a **dns_name** is set, TOSCA Orchestrators SHALL attempt to register the name in the (local)
2754 DNS registry for the Cloud provider.

2755 **5.5.9 tosca.capabilities.Endpoint.Admin**

2756 This is the default TOSCA type that should be used or extended to define a specialized administrator
2757 endpoint capability.

Shorthand Name	Endpoint.Admin
Type Qualified Name	tosca:Endpoint.Admin
Type URI	tosca.capabilities.Endpoint.Admin

2758 **5.5.9.1 Properties**

Name	Required	Type	Constraints	Description
None	N/A	N/A	N/A	N/A

2759 **5.5.9.2 Definition**

```
tosca.capabilities.Endpoint.Admin:
  derived_from: tosca.capabilities.Endpoint
  # Change Endpoint secure indicator to true from its default of false
  properties:
    secure:
      type: boolean
      default: true
      constraints:
        - equal: true
```

2760 **5.5.9.3 Additional requirements**

- 2761 • TOSCA Orchestrator implementations of Endpoint.Admin (and connections to it) **SHALL** assure
2762 that network-level security is enforced if possible.

2763 **5.5.10 tosca.capabilities.Endpoint.Database**

2764 This is the default TOSCA type that should be used or extended to define a specialized database
2765 endpoint capability.

Shorthand Name	Endpoint.Database
Type Qualified Name	tosca:Endpoint.Database
Type URI	tosca.capabilities.Endpoint.Database

2766 **5.5.10.1 Properties**

Name	Required	Type	Constraints	Description
None	N/A	N/A	N/A	N/A

2767 **5.5.10.2 Definition**

```
tosca.capabilities.Endpoint.Database:
  derived_from: tosca.capabilities.Endpoint
```

2768 **5.5.11 tosca.capabilities.Attachment**

2769 This is the default TOSCA type that should be used or extended to define an attachment capability of a
2770 (logical) infrastructure device node (e.g., [BlockStorage](#) node).

Shorthand Name	Attachment
Type Qualified Name	tosca:Attachment
Type URI	tosca.capabilities.Attachment

2771 **5.5.11.1 Properties**

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

2772 **5.5.11.2 Definition**

```
tosca.capabilities.Attachment:
  derived_from: tosca.capabilities.Root
```

2773 **5.5.12 tosca.capabilities.OperatingSystem**

2774 This is the default TOSCA type that should be used to express an Operating System capability for a
2775 node.

Shorthand Name	OperatingSystem
Type Qualified Name	tosca:OperatingSystem
Type URI	tosca.capabilities.OperatingSystem

2776 **5.5.12.1 Properties**

Name	Required	Type	Constraints	Description
architecture	no	string	None	The Operating System (OS) architecture. Examples of valid values include: x86_32, x86_64, etc.
type	no	string	None	The Operating System (OS) type. Examples of valid values include: linux, aix, mac, windows, etc.
distribution	no	string	None	The Operating System (OS) distribution. Examples of valid values for an "type" of "Linux" would include: debian, fedora, rhel and ubuntu.
version	no	version	None	The Operating System version.

2777 **5.5.12.2 Definition**

```
tosca.capabilities.OperatingSystem:
  derived_from: tosca.capabilities.Root
  properties:
    architecture:
      type: string
      required: false
    type:
      type: string
      required: false
    distribution:
      type: string
      required: false
    version:
      type: version
      required: false
```

2778 **5.5.12.3 Additional Requirements**

- 2779 • Please note that the string values for the properties **architecture**, **type** and **distribution**
2780 SHALL be normalized to lowercase by processors of the service template for matching purposes.

2781 For example, if a “**type**” value is set to either “Linux”, “LINUX” or “linux” in a service template, the
 2782 processor would normalize all three values to “linux” for matching purposes.

2783 **5.5.13 tosca.capabilities.Scalable**

2784 This is the default TOSCA type that should be used to express a scalability capability for a node.

Shorthand Name	Scalable
Type Qualified Name	tosca:Scalable
Type URI	tosca.capabilities.Scalable

2785 **5.5.13.1 Properties**

Name	Required	Type	Constraints	Description
min_instances	yes	integer	default: 1	This property is used to indicate the minimum number of instances that should be created for the associated TOSCA Node Template by a TOSCA orchestrator.
max_instances	yes	integer	default: 1	This property is used to indicate the maximum number of instances that should be created for the associated TOSCA Node Template by a TOSCA orchestrator.
default_instances	no	integer	N/A	An optional property that indicates the requested default number of instances that should be the starting number of instances a TOSCA orchestrator should attempt to allocate. Note: The value for this property MUST be in the range between the values set for ‘min_instances’ and ‘max_instances’ properties.

2786 **5.5.13.2 Definition**

```
tosca.capabilities.Scalable:
  derived_from: tosca.capabilities.Root
  properties:
    min_instances:
      type: integer
      default: 1
    max_instances:
      type: integer
      default: 1
    default_instances:
      type: integer
```

2787 **5.5.13.3 Notes**

- 2788 • The actual number of instances for a node may be governed by a separate scaling policy which
 2789 conceptually would be associated to either a scaling-capable node or a group of nodes in which it
 2790 is defined to be a part of. This is a planned future feature of the TOSCA Simple Profile and not
 2791 currently described.

2792 **5.5.14 tosca.capabilities.network.Bindable**

2793 A node type that includes the Bindable capability indicates that it can be bound to a logical network
2794 association via a network port.

Shorthand Name	network.Bindable
Type Qualified Name	tosca:network.Bindable
Type URI	tosca.capabilities.network.Bindable

2795 **5.5.14.1 Properties**

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

2796 **5.5.14.2 Definition**

```
tosca.capabilities.network.Bindable:
  derived_from: tosca.capabilities.Node
```

2797 **5.6 Requirement Types**

2798 There are no normative Requirement Types currently defined in this working draft. Typically,
2799 Requirements are described against a known Capability Type

2800 **5.7 Relationship Types**

2801 **5.7.1 tosca.relationships.Root**

2802 This is the default (root) TOSCA Relationship Type definition that all other TOSCA Relationship Types
2803 derive from.

2804 **5.7.1.1 Attributes**

Name	Required	Type	Constraints	Description
tosca_id	yes	string	None	A unique identifier of the realized instance of a Relationship Template that derives from any TOSCA normative type.
tosca_name	yes	string	None	This attribute reflects the name of the Relationship Template as defined in the TOSCA service template. This name is not unique to the realized instance model of corresponding deployed application as each template in the model can result in one or more instances (e.g., scaled) when orchestrated to a provider environment.
state	yes	string	default: initial	The state of the relationship instance. See section "Relationship States" for allowed values.

2805 **5.7.1.2 Definition**

```
tosca.relationships.Root:
  description: The TOSCA root Relationship Type all other TOSCA base Relationship
  Types derive from
  attributes:
    tosca_id:
```

```
    type: string
  tosca_name:
    type: string
  interfaces:
    Configure:
      type: tosca.interfaces.relationship.Configure
```

2806 **5.7.2 tosca.relationships.DependsOn**

2807 This type represents a general dependency relationship between two nodes.

Shorthand Name	DependsOn
Type Qualified Name	tosca:DependsOn
Type URI	tosca.relationships.DependsOn

2808 **5.7.2.1 Definition**

```
tosca.relationships.DependsOn:
  derived_from: tosca.relationships.Root
  valid_target_types: [ tosca.capabilities.Node ]
```

2809 **5.7.3 tosca.relationships.HostedOn**

2810 This type represents a hosting relationship between two nodes.

Shorthand Name	HostedOn
Type Qualified Name	tosca:HostedOn
Type URI	tosca.relationships.HostedOn

2811 **5.7.3.1 Definition**

```
tosca.relationships.HostedOn:
  derived_from: tosca.relationships.Root
  valid_target_types: [ tosca.capabilities.Container ]
```

2812 **5.7.4 tosca.relationships.ConnectsTo**

2813 This type represents a network connection relationship between two nodes.

Shorthand Name	ConnectsTo
Type Qualified Name	tosca:ConnectsTo
Type URI	tosca.relationships.ConnectsTo

2814 **5.7.4.1 Definition**

```
tosca.relationships.ConnectsTo:
  derived_from: tosca.relationships.Root
  valid_target_types: [ tosca.capabilities.Endpoint ]
  properties:
    credential:
      type: tosca.datatypes.Credential
      required: false
```


2815 **5.7.4.2 Properties**

Name	Required	Type	Constraints	Description
credential	no	Credential	None	The security credential to use to present to the target endpoint to for either authentication or authorization purposes.

2816 **5.7.5 tosca.relationships.AttachesTo**

2817 This type represents an attachment relationship between two nodes. For example, an AttachesTo
 2818 relationship type would be used for attaching a storage node to a Compute node.

Shorthand Name	AttachesTo
Type Qualified Name	tosca:AttachesTo
Type URI	tosca.relationships.AttachesTo

2819 **5.7.5.1 Properties**

Name	Required	Type	Constraints	Description
location	yes	string	min_length: 1	The relative location (e.g., path on the file system), which provides the root location to address an attached node. e.g., a mount point / path such as '/usr/data' Note: The user must provide it and it cannot be "root".
device	no	string	None	The logical device name which for the attached device (which is represented by the target node in the model). e.g., '/dev/hda1'

2820 **5.7.5.2 Attributes**

Name	Required	Type	Constraints	Description
device	no	string	None	The logical name of the device as exposed to the instance. Note: A runtime property that gets set when the model gets instantiated by the orchestrator.

2821 **5.7.5.3 Definition**

```
tosca.relationships.AttachesTo:
  derived_from: tosca.relationships.Root
  valid_target_types: [ tosca.capabilities.Attachment ]
  properties:
    location:
      type: string
      constraints:
        - min_length: 1
    device:
      type: string
      required: false
```

2822 **5.7.6 tosca.relationships.RoutesTo**

2823 This type represents an intentional network routing between two Endpoints in different networks.

Shorthand Name	RoutesTo
Type Qualified Name	tosca:RoutesTo
Type URI	tosca.relationships.RoutesTo

2824 **5.7.6.1 Definition**

```
tosca.relationships.RoutesTo:  
  derived_from: tosca.relationships.ConnectsTo  
  valid_target_types: [ tosca.capabilities.Endpoint ]
```

2825 **5.8 Interface Types**

2826 Interfaces are reusable entities that define a set of operations that that can be included as part of a Node
2827 type or Relationship Type definition. Each named operations may have code or scripts associated with
2828 them that orchestrators can execute for when transitioning an application to a given state.

2829 **5.8.1 Additional Requirements**

- 2830
- 2831 • Designers of Node or Relationship types are not required to actually provide/associate code or
2832 scripts with every operation for a given interface it supports. In these cases, orchestrators SHALL
2833 consider that a “No Operation” or “no-op”.
 - 2834 • The default behavior when providing scripts for an operation in a sub-type (sub-class) or a
2835 template of an existing type which already has a script provided for that operation SHALL be
override. Meaning that the subclasses’ script is used in place of the parent type’s script.

2836 **5.8.2 Best Practices**

- 2837 • When TOSCA Orchestrators substitute an implementation for an abstract node in a deployed
2838 service template it SHOULD be able to present a confirmation to the submitter to confirm the
2839 implementation chosen would be acceptable.

2840 **5.8.3 tosca.interfaces.Root**

2841 This is the default (root) TOSCA Interface Type definition that all other TOSCA Interface Types derive
2842 from.

2843 **5.8.3.1 Definition**

```
tosca.interfaces.Root:  
  derived_from: tosca.entity.Root  
  description: The TOSCA root Interface Type all other TOSCA base Interface Types  
  derive from
```

2844 **5.8.4 tosca.interfaces.node.lifecycle.Standard**

2845 This lifecycle interface defines the essential, normative operations that TOSCA nodes may support.

Shorthand Name	Standard
Type Qualified Name	tosca:Standard
Type URI	tosca.interfaces.node.lifecycle.Standard

2846 **5.8.4.1 Definition**

```
tosca.interfaces.node.lifecycle.Standard:
  derived_from: toska.interfaces.Root
  create:
    description: Standard lifecycle create operation.
  configure:
    description: Standard lifecycle configure operation.
  start:
    description: Standard lifecycle start operation.
  stop:
    description: Standard lifecycle stop operation.
  delete:
    description: Standard lifecycle delete operation.
```

2847 **5.8.4.2 Create operation**

2848 The create operation is generally used to create the resource or service the node represents in the
 2849 topology. TOSCA orchestrators expect node templates to provide either a deployment artifact or an
 2850 implementation artifact of a defined artifact type that it is able to process. This specification defines
 2851 normative deployment and implementation artifact types all TOSCA Orchestrators are expected to be
 2852 able to process to support application portability.

2853 **5.8.4.3 TOSCA Orchestrator processing of Deployment artifacts**

2854 TOSCA Orchestrators, when encountering a deployment artifact on the create operation; will
 2855 automatically attempt to deploy the artifact based upon its artifact type. This means that no
 2856 implementation artifacts (e.g., scripts) are needed on the create operation to provide commands that
 2857 deploy or install the software.

2858
 2859 For example, if a TOSCA Orchestrator is processing an application with a node of type
 2860 SoftwareComponent and finds that the node's template has a create operation that provides a filename
 2861 (or references to an artifact which describes a file) of a known TOSCA deployment artifact type such as
 2862 an Open Virtualization Format (OVF) image it will automatically deploy that image into the
 2863 SoftwareComponent's host Compute node.

2864 **5.8.4.4 Operation sequencing and node state**

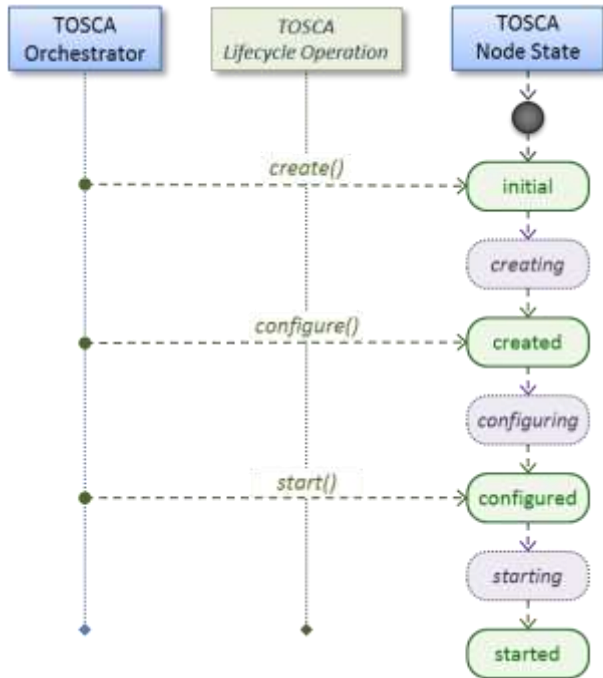
2865 The following diagrams show how TOSCA orchestrators sequence the operations of the Standard
 2866 lifecycle in normal node startup and shutdown procedures.

2867 The following key should be used to interpret the diagrams:



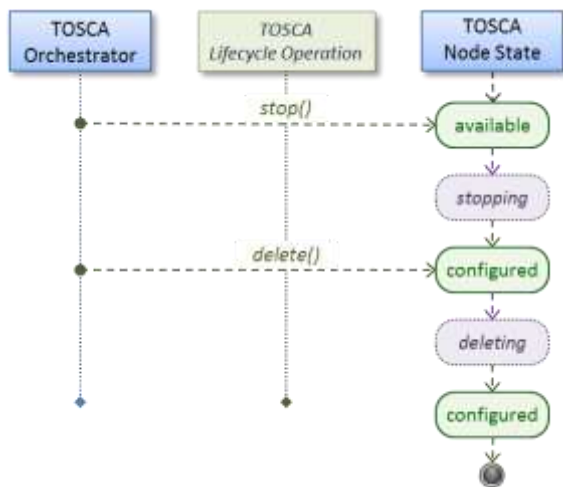
2868 **5.8.4.4.1 Normal node startup sequence diagram**

2869 The following diagram shows how the TOSCA orchestrator would invoke operations on the Standard
 2870 lifecycle to startup a node.



2871 **5.8.4.4.2 Normal node shutdown sequence diagram**

2872 The following diagram shows how the TOSCA orchestrator would invoke operations on the Standard
 2873 lifecycle to shut down a node.



2874

2875 **5.8.5 tosca.interfaces.relationship.Configure**

2876 The lifecycle interfaces define the essential, normative operations that each TOSCA Relationship Types
2877 may support.

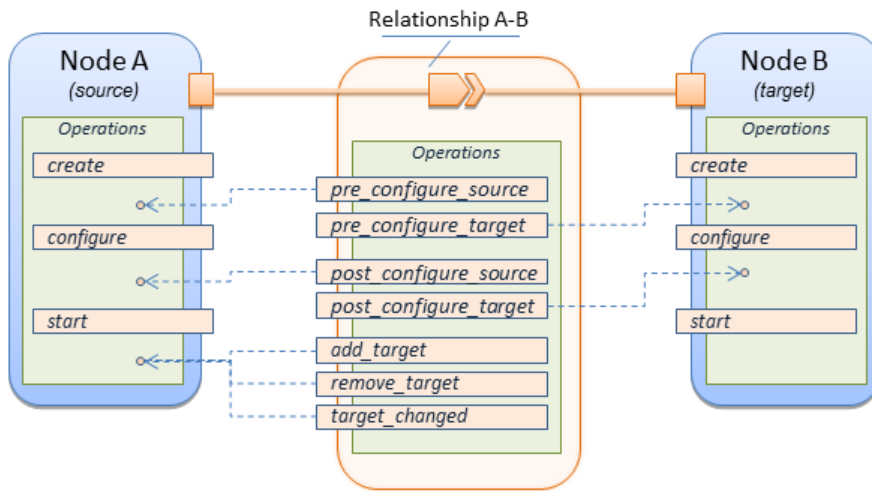
Shorthand Name	Configure
Type Qualified Name	tosca:Configure
Type URI	tosca.interfaces.relationship.Configure

2878 **5.8.5.1 Definition**

```
tosca.interfaces.relationship.Configure:  
  derived_from: tosca.interfaces.Root  
  pre_configure_source:  
    description: Operation to pre-configure the source endpoint.  
  pre_configure_target:  
    description: Operation to pre-configure the target endpoint.  
  post_configure_source:  
    description: Operation to post-configure the source endpoint.  
  post_configure_target:  
    description: Operation to post-configure the target endpoint.  
  add_target:  
    description: Operation to notify the source node of a target node being added  
via a relationship.  
  add_source:  
    description: Operation to notify the target node of a source node which is now  
available via a relationship.  
  description:  
  target_changed:  
    description: Operation to notify source some property or attribute of the  
target changed  
  remove_target:  
    description: Operation to remove a target node.
```

2879

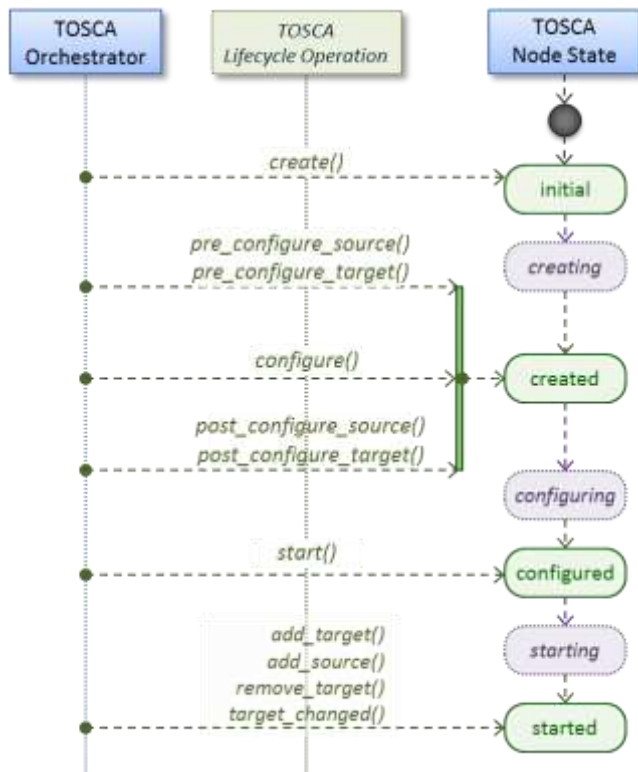
2880 **5.8.5.2 Invocation Conventions**



2881 TOSCA relationships are directional connecting a source node to a target node. When TOSCA
 2882 Orchestrator connects a source and target node together using a relationship that supports the Configure
 2883 interface it will “interleave” the operations invocations of the Configure interface with those of the node’s
 2884 own Standard lifecycle interface. This concept is illustrated below:

2885 **5.8.5.3 Normal node start sequence with Configure relationship operations**

2886 The following diagram shows how the TOSCA orchestrator would invoke Configure lifecycle operations in
 2887 conjunction with Standard lifecycle operations during a typical startup sequence on a node.



2888 **5.8.5.4 Node-Relationship configuration sequence**

2889 Depending on which side (i.e., source or target) of a relationship a node is on, the orchestrator will:

- 2890 • Invoke either the **pre_configure_source** or **pre_configure_target** operation as supplied by
- 2891 the relationship on the node.
- 2892 • Invoke the node's **configure** operation.
- 2893 • Invoke either the **post_configure_source** or **post_configure_target** as supplied by the
- 2894 relationship on the node.

2895 Note that the **pre_configure_xxx** and **post_configure_xxx** are invoked only once per node instance.

2896 **5.8.5.4.1 Node-Relationship add, remove and changed sequence**

2897 Since a topology template contains nodes that can dynamically be added (and scaled), removed or

2898 changed as part of an application instance, the Configure lifecycle includes operations that are invoked

2899 on node instances that to notify and address these dynamic changes.

2900

2901 For example, a source node, of a relationship that uses the Configure lifecycle, will have the relationship

2902 operations **add_target**, or **remove_target** invoked on it whenever a target node instance is added or

2903 removed to the running application instance. In addition, whenever the node state of its target node

2904 changes, the **target_changed** operation is invoked on it to address this change. Conversely, the

2905 **add_source** and **remove_source** operations are invoked on the source node of the relationship.

2906 **5.8.5.5 Notes**

- 2907 • The target (provider) **MUST** be active and running (i.e., all its dependency stack **MUST** be
- 2908 fulfilled) prior to invoking **add_target**
- 2909 • In other words, all Requirements **MUST** be satisfied before it advertises its capabilities (i.e.,
- 2910 the attributes of the matched Capabilities are available).
- 2911 • In other words, it cannot be “consumed” by any dependent node.
- 2912 • Conversely, since the source (consumer) needs information (attributes) about any targets
- 2913 (and their attributes) being removed before it actually goes away.
- 2914 • The **remove_target** operation should only be executed if the target has had **add_target**
- 2915 executed. **BUT** in truth we're first informed about a target in **pre_configure_source**, so if we
- 2916 execute that the source node should see **remove_target** called to cleanup.
- 2917 • **Error handling:** If any node operation of the topology fails processing should stop on that node
- 2918 template and the failing operation (script) should return an error (failure) code when possible.

2919 **5.9 Node Types**

2920 **5.9.1 tosca.nodes.Root**

2921 The TOSCA **Root** Node Type is the default type that all other TOSCA base Node Types derive from.

2922 This allows for all TOSCA nodes to have a consistent set of features for modeling and management (e.g.,

2923 consistent definitions for requirements, capabilities and lifecycle interfaces).

2924

Shorthand Name	Root
Type Qualified Name	tosca:Root
Type URI	tosca.nodes.Root

2925 **5.9.1.1 Properties**

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	The TOSCA Root Node type has no specified properties.

2926 **5.9.1.2 Attributes**

Name	Required	Type	Constraints	Description
tosca_id	yes	string	None	A unique identifier of the realized instance of a Node Template that derives from any TOSCA normative type.
tosca_name	yes	string	None	This attribute reflects the name of the Node Template as defined in the TOSCA service template. This name is not unique to the realized instance model of corresponding deployed application as each template in the model can result in one or more instances (e.g., scaled) when orchestrated to a provider environment.
state	yes	string	default: initial	The state of the node instance. See section “Node States” for allowed values.

2927 **5.9.1.3 Definition**

```

tosca.nodes.Root:
  derived_from: tosca.entity.Root
  description: The TOSCA Node Type all other TOSCA base Node Types derive from
  attributes:
    tosca_id:
      type: string
    tosca_name:
      type: string
    state:
      type: string
  capabilities:
    feature:
      type: tosca.capabilities.Node
  requirements:
    - dependency:
        capability: tosca.capabilities.Node
        node: tosca.nodes.Root
        relationship: tosca.relationships.DependsOn
        occurrences: [ 0, UNBOUNDED ]
  interfaces:
    Standard:
      type: tosca.interfaces.node.lifecycle.Standard

```

2928 **5.9.1.4 Additional Requirements**

- 2929 • All Node Type definitions that wish to adhere to the TOSCA Simple Profile **SHOULD** extend from the
2930 TOSCA Root Node Type to be assured of compatibility and portability across implementations.

2931 **5.9.2 tosca.nodes.Compute**

2932 The TOSCA **Compute** node represents one or more real or virtual processors of software applications or
 2933 services along with other essential local resources. Collectively, the resources the compute node
 2934 represents can logically be viewed as a (real or virtual) “server”.

Shorthand Name	Compute
Type Qualified Name	tosca:Compute
Type URI	tosca.nodes.Compute

2935 **5.9.2.1 Properties**

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

2936 **5.9.2.2 Attributes**

Name	Required	Type	Constraints	Description
private_address	no	string	None	The primary private IP address assigned by the cloud provider that applications may use to access the Compute node.
public_address	no	string	None	The primary public IP address assigned by the cloud provider that applications may use to access the Compute node.
networks	no	map of NetworkInfo	None	The list of logical networks assigned to the compute host instance and information about them.
ports	no	map of PortInfo	None	The list of logical ports assigned to the compute host instance and information about them.

2937 **5.9.2.3 Definition**

```

tosca.nodes.Compute:
  derived_from: tosca.nodes.Root
  attributes:
    private_address:
      type: string
    public_address:
      type: string
    networks:
      type: map
      entry_schema:
        type: tosca.datatypes.network.NetworkInfo
    ports:
      type: map
      entry_schema:
        type: tosca.datatypes.network.PortInfo
  requirements:
    - local_storage:
      capability: tosca.capabilities.Attachment
      node: tosca.nodes.BlockStorage
      relationship: tosca.relationships.AttachesTo
  
```

```

    occurrences: [0, UNBOUNDED]
capabilities:
  host:
    type: tosca.capabilities.Container
    valid_source_types: [tosca.nodes.SoftwareComponent]
  endpoint:
    type: tosca.capabilities.Endpoint.Admin
  os:
    type: tosca.capabilities.OperatingSystem
  scalable:
    type: tosca.capabilities.Scalable
  binding:
    type: tosca.capabilities.network.Bindable

```

2938 **5.9.2.4 Additional Requirements**

- 2939
- The underlying implementation of the Compute node SHOULD have the ability to instantiate guest operating systems (either actual or virtualized) based upon the OperatingSystem capability properties if they are supplied in the a node template derived from the Compute node type.
- 2940
- 2941

2942 **5.9.3 tosca.nodes.SoftwareComponent**

2943 The TOSCA **SoftwareComponent** node represents a generic software component that can be managed
 2944 and run by a TOSCA **Compute** Node Type.

Shorthand Name	SoftwareComponent
Type Qualified Name	tosca:SoftwareComponent
Type URI	tosca.nodes.SoftwareComponent

2945 **5.9.3.1 Properties**

Name	Required	Type	Constraints	Description
component_version	no	version	None	The optional software component's version.
admin_credential	no	Credential	None	The optional credential that can be used to authenticate to the software component.

2946 **5.9.3.2 Attributes**

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

2947 **5.9.3.3 Definition**

```

tosca.nodes.SoftwareComponent:
  derived_from: tosca.nodes.Root
  properties:
    # domain-specific software component version
    component_version:
      type: version
      required: false
    admin_credential:
      type: tosca.datatypes.Credential

```

```

    required: false
  requirements:
    - host:
        capability: tosca.capabilities.Container
        node: tosca.nodes.Compute
        relationship: tosca.relationships.HostedOn

```

2948 5.9.3.4 Additional Requirements

- 2949
- Nodes that can directly be managed and run by a TOSCA **Compute** Node Type **SHOULD** extend from this type.
- 2950

2951 5.9.4 [tosca.nodes.WebServer](#)

2952 This TOSA **WebServer** Node Type represents an abstract software component or service that is capable
 2953 of hosting and providing management operations for one or more **WebApplication** nodes.

Shorthand Name	WebServer
Type Qualified Name	tosca:WebServer
Type URI	tosca.nodes.WebServer

2954 5.9.4.1 Properties

Name	Required	Type	Constraints	Description
None	N/A	N/A	N/A	N/A

2955 5.9.4.2 Definition

```

tosca.nodes.WebServer:
  derived_from: tosca.nodes.SoftwareComponent
  capabilities:
    # Private, layer 4 endpoints
    data_endpoint: tosca.capabilities.Endpoint
    admin_endpoint: tosca.capabilities.Endpoint.Admin
  host:
    type: tosca.capabilities.Container
    valid_source_types: [ tosca.nodes.WebApplication ]

```

2956 5.9.4.3 Additional Requirements

- 2957
- This node **SHALL** export both a secure endpoint capability (i.e., **admin_endpoint**), typically for
 2958 administration, as well as a regular endpoint (i.e., **data_endpoint**) for serving data.

2959 5.9.5 [tosca.nodes.WebApplication](#)

2960 The TOSCA **WebApplication** node represents a software application that can be managed and run by a
 2961 TOSCA **WebServer** node. Specific types of web applications such as Java, etc. could be derived from
 2962 this type.

Shorthand Name	WebApplication
Type Qualified Name	tosca: WebApplication
Type URI	tosca.nodes.WebApplication

2963 **5.9.5.1 Properties**

Name	Required	Type	Constraints	Description
context_root	no	string	None	The web application's context root which designates the application's URL path within the web server it is hosted on.

2964 **5.9.5.2 Definition**

```

tosca.nodes.WebApplication:
  derived_from: tosca.nodes.Root
  properties:
    context_root:
      type: string
  capabilities:
    app_endpoint:
      type: tosca.capabilities.Endpoint
  requirements:
    - host:
      capability: tosca.capabilities.Container
      node: tosca.nodes.WebServer
      relationship: tosca.relationships.HostedOn

```

2965 **5.9.6 toasca.nodes.DBMS**

2966 The TOSCA DBMS node represents a typical relational, SQL Database Management System software
 2967 component or service.

2968 **5.9.6.1 Properties**

Name	Required	Type	Constraints	Description
root_password	no	string	None	The optional root password for the DBMS server.
port	no	integer	None	The DBMS server's port.

2969 **5.9.6.2 Definition**

```

tosca.nodes.DBMS:
  derived_from: tosca.nodes.SoftwareComponent
  properties:
    root_password:
      type: string
      required: false
      description: the optional root password for the DBMS service
    port:
      type: integer

```

```

required: false
description: the port the DBMS service will listen to for data and requests
capabilities:
  host:
    type: tosca.capabilities.Container
    valid_source_types: [ tosca.nodes.Database ]

```

2970 **5.9.7 tosca.nodes.Database**

2971 The TOSCA **Database** node represents a logical database that can be managed and hosted by a TOSCA
 2972 **DBMS** node.

Shorthand Name	Database
Type Qualified Name	tosca:Database
Type URI	tosca.nodes.Database

2973 **5.9.7.1 Properties**

Name	Required	Type	Constraints	Description
name	yes	string	None	The logical database Name
port	no	integer	None	The port the database service will use to listen for incoming data and requests.
user	no	string	None	The special user account used for database administration.
password	no	string	None	The password associated with the user account provided in the 'user' property.

2974 **5.9.7.2 Definition**

```

tosca.nodes.Database:
  derived_from: tosca.nodes.Root
  properties:
    name:
      type: string
      description: the logical name of the database
    port:
      type: integer
      description: the port the underlying database service will listen to for data
    user:
      type: string
      description: the optional user account name for DB administration
      required: false
    password:
      type: string
      description: the optional password for the DB user account
      required: false
  requirements:
    - host:
        capability: tosca.capabilities.Container
        node: tosca.nodes.DBMS
        relationship: tosca.relationships.HostedOn
  capabilities:
    database_endpoint:
      type: tosca.capabilities.Endpoint.Database

```

2975 **5.9.8 tosca.nodes.Storage.ObjectStorage**

2976 The TOSCA **ObjectStorage** node represents storage that provides the ability to store data as objects (or
 2977 BLOBs of data) without consideration for the underlying filesystem or devices.

Shorthand Name	ObjectStorage
Type Qualified Name	tosca:ObjectStorage
Type URI	tosca.nodes.Storage.ObjectStorage

2978 **5.9.8.1 Properties**

Name	Required	Type	Constraints	Description
name	yes	string	None	The logical name of the object store (or container).
size	no	scalar-unit.size	greater_or_equal: 0 GB	The requested initial storage size (default unit is in Gigabytes).
maxsize	no	scalar-unit.size	greater_or_equal: 0 GB	The requested maximum storage size (default unit is in Gigabytes).

2979 **5.9.8.2 Definition**

```

tosca.nodes.Storage.ObjectStorage:
  derived_from: tosca.nodes.Root
  properties:
    name:
      type: string
    size:
      type: scalar-unit.size
      constraints:
        - greater_or_equal: 0 GB
    maxsize:
      type: scalar-unit.size
      constraints:
        - greater_or_equal: 0 GB
  capabilities:
    storage_endpoint:
      type: tosca.capabilities.Endpoint

```

2980 **5.9.8.3 Notes:**

- 2981 • Subclasses of the `tosca.nodes.ObjectStorage` node type may impose further constraints on
 2982 properties. For example, a subclass may constrain the (minimum or maximum) length of the
 2983 'name' property or include a regular expression to constrain allowed characters used in the
 2984 'name' property.

2985 **5.9.9 tosca.nodes.Storage.BlockStorage**

2986 The TOSCA **BlockStorage** node currently represents a server-local block storage device (i.e., not
 2987 shared) offering evenly sized blocks of data from which raw storage volumes can be created.

2988 **Note:** In this draft of the TOSCA Simple Profile, distributed or Network Attached Storage (NAS) are not
 2989 yet considered (nor are clustered file systems), but the TC plans to do so in future drafts.

Shorthand Name	BlockStorage
Type Qualified Name	tosca:BlockStorage
Type URI	tosca.nodes.Storage.BlockStorage

2990 **5.9.9.1 Properties**

Name	Required	Type	Constraints	Description
size	yes *	scalar-unit.size	greater_or_equal: 1 MB	The requested storage size (default unit is MB). * Note: <ul style="list-style-type: none"> • Required when an existing volume (i.e., <code>volume_id</code>) is not available. • If <code>volume_id</code> is provided, size is ignored. Resize of existing volumes is not considered at this time.
volume_id	no	string	None	ID of an existing volume (that is in the accessible scope of the requesting application).

Name	Required	Type	Constraints	Description
snapshot_id	no	string	None	Some identifier that represents an existing snapshot that should be used when creating the block storage (volume).

2991 **5.9.9.2 Attributes**

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

2992 **5.9.9.3 Definition**

```

tosca.nodes.Storage.BlockStorage:
  derived_from: tosca.nodes.Root
  properties:
    size:
      type: scalar-unit.size
      constraints:
        - greater_or_equal: 1 MB
    volume_id:
      type: string
      required: false
    snapshot_id:
      type: string
      required: false
  capabilities:
    attachment:
      type: tosca.capabilities.Attachment

```

2993 **5.9.9.4 Additional Requirements**

- 2994
- The **size** property is required when an existing volume (i.e., **volume_id**) is not available.
- 2995
- However, if the property **volume_id** is provided, the **size** property is ignored.

2996 **5.9.9.5 Notes**

- 2997
- Resize of existing volumes is not considered at this time.
- 2998
- It is assumed that the volume contains a single filesystem that the operating system (that is hosting an associate application) can recognize and mount without additional information (i.e., it is operating system independent).
- 2999
- Currently, this version of the Simple Profile does not consider regions (or availability zones) when modeling storage.
- 3000
- 3001
- 3002

3003 **5.9.10 tosca.nodes.Container.Runtime**

3004 The TOSCA **Container** Runtime node represents operating system-level virtualization technology used

3005 to run multiple application services on a single Compute host.

Shorthand Name	Container.Runtime
Type Qualified Name	tosca:Container.Runtime
Type URI	tosca.nodes.Container.Runtime

3006 **5.9.10.1 Definition**

```
tosca.nodes.Container.Runtime:
  derived_from: tosca.nodes.SoftwareComponent
  capabilities:
    host:
      type: tosca.capabilities.Container
  scalable:
    type: tosca.capabilities.Scalable
```

3007 **5.9.11 tosca.nodes.Container.Application**

3008 The TOSCA **Container Application** node represents an application that requires **Container**-level
3009 virtualization technology.

Shorthand Name	Container.Application
Type Qualified Name	tosca:Container.Application
Type URI	tosca.nodes.Container.Application

3010 **5.9.11.1 Definition**

```
tosca.nodes.Container.Application:
  derived_from: tosca.nodes.Root
  requirements:
    - host:
      capability: tosca.capabilities.Container
      node: tosca.nodes.Container.Runtime
      relationship: tosca.relationships.HostedOn
    - storage:
      capability: tosca.capabilities.Storage
    - network:
      capability: tosca.capabilities.EndPoint
```

3011 **5.9.12 tosca.nodes.LoadBalancer**

3012 The TOSCA **Load Balancer** node represents logical function that be used in conjunction with a Floating
3013 Address to distribute an application's traffic (load) across a number of instances of the application (e.g.,
3014 for a clustered or scaled application).

Shorthand Name	LoadBalancer
Type Qualified Name	tosca:LoadBalancer
Type URI	tosca.nodes.LoadBalancer

3015 **5.9.12.1 Definition**

```
tosca.nodes.LoadBalancer:
```

```

derived_from: tosca.nodes.Root
properties:
  algorithm:
    type: string
    required: false
    status: experimental
capabilities:
  client:
    type: tosca.capabilities.Endpoint.Public
    occurrences: [0, UNBOUNDED]
    description: the Floating (IP) client's on the public network can connect to
requirements:
  - application:
    capability: tosca.capabilities.Endpoint
    relationship: tosca.relationships.RoutesTo
    occurrences: [0, UNBOUNDED]
    description: Connection to one or more load balanced applications

```

3016 5.9.12.2 Notes:

- 3017
- A **LoadBalancer** node can still be instantiated and managed independently of any applications it would serve; therefore, the load balancer's **application** requirement allows for zero
- 3018
- 3019

3020 5.10 Group Types

3021 TOSCA Group Types represent logical groupings of TOSCA nodes that have an implied membership
 3022 relationship and may need to be orchestrated or managed together to achieve some result. Some use
 3023 cases being developed by the TOSCA TC use groups to apply TOSCA policies for software placement
 3024 and scaling while other use cases show groups can be used to describe cluster relationships.

3025

3026 **Note:** Additional normative TOSCA Group Types and use cases for them will be developed in future
 3027 drafts of this specification.

3028 5.10.1 [tosca.groups.Root](#)

3029 This is the default (root) TOSCA [Group Type](#) definition that all other TOSCA base Group Types derive
 3030 from.

3031 5.10.1.1 Definition

```

tosca.groups.Root:
  description: The TOSCA Group Type all other TOSCA Group Types derive from
  interfaces:
    Standard:
      type: tosca.interfaces.node.lifecycle.Standard

```

3032 5.10.1.2 Notes:

- 3033
- Group operations are not necessarily tied directly to member nodes that are part of a group.
- 3034
- Future versions of this specification will create sub types of the **tosca.groups.Root** type that will
- 3035

3036 **5.11 Policy Types**

3037 TOSCA Policy Types represent logical grouping of TOSCA nodes that have an implied relationship and
3038 need to be orchestrated or managed together to achieve some result. Some use cases being developed
3039 by the TOSCA TC use groups to apply TOSCA policies for software placement and scaling while other
3040 use cases show groups can be used to describe cluster relationships.

3041 **5.11.1 `tosca.policies.Root`**

3042 This is the default (root) TOSCA Policy Type definition that all other TOSCA base Policy Types derive
3043 from.

3044 **5.11.1.1 Definition**

```
tosca.policies.Root:  
  description: The TOSCA Policy Type all other TOSCA Policy Types derive from
```

3045 **5.11.2 `tosca.policies.Placement`**

3046 This is the default (root) TOSCA Policy Type definition that is used to govern placement of TOSCA nodes
3047 or groups of nodes.

3048 **5.11.2.1 Definition**

```
tosca.policies.Placement:  
  derived_from: tosca.policies.Root  
  description: The TOSCA Policy Type definition that is used to govern placement of  
  TOSCA nodes or groups of nodes.
```

3049 **5.11.3 `tosca.policies.Scaling`**

3050 This is the default (root) TOSCA Policy Type definition that is used to govern scaling of TOSCA nodes or
3051 groups of nodes.

3052 **5.11.3.1 Definition**

```
tosca.policies.Scaling:  
  derived_from: tosca.policies.Root  
  description: The TOSCA Policy Type definition that is used to govern scaling of  
  TOSCA nodes or groups of nodes.
```

3053 **5.11.4 `tosca.policies.Update`**

3054 This is the default (root) TOSCA Policy Type definition that is used to govern update of TOSCA nodes or
3055 groups of nodes.

3056 **5.11.4.1 Definition**

```
tosca.policies.Update:  
  derived_from: tosca.policies.Root  
  description: The TOSCA Policy Type definition that is used to govern update of  
  TOSCA nodes or groups of nodes.
```

3057 **5.11.5 tosca.policies.Performance**

3058 This is the default (root) TOSCA Policy Type definition that is used to declare performance requirements
3059 for TOSCA nodes or groups of nodes.

3060 **5.11.5.1 Definition**

```
tosca.policies.Performance:  
  derived_from: tosca.policies.Root  
  description: The TOSCA Policy Type definition that is used to declare performance  
  requirements for TOSCA nodes or groups of nodes.
```

3061

3062 6 TOSCA Cloud Service Archive (CSAR) format

3063 Except for the examples, this section is **normative** and defines changes to the TOSCA archive format
3064 relative to the TOSCA v1.0 XML specification.

3065
3066 TOSCA Simple Profile definitions along with all accompanying artifacts (e.g. scripts, binaries,
3067 configuration files) can be packaged together in a CSAR file as already defined in the TOSCA version 1.0
3068 specification [**TOSCA-1.0**]. In contrast to the TOSCA 1.0 CSAR file specification (see chapter 16 in
3069 [**TOSCA-1.0**]), this simple profile makes a few simplifications both in terms of overall CSAR file structure
3070 as well as meta-file content as described below.

3071 6.1 Overall Structure of a CSAR

3072 A CSAR zip file is required to contain one of the following:

- 3073 • a **TOSCA-Metadata** directory, which in turn contains the **TOSCA.meta** metadata file that provides
3074 entry information for a TOSCA orchestrator processing the CSAR file.
- 3075 • a yaml (.yml or .yaml) file at the root of the archive. The yaml file being a valid toasca definition
3076 template that **MUST** define a metadata section where `template_name` and `template_version` are
3077 required.

3078 The CSAR file may contain other directories with arbitrary names and contents. Note that in contrast to
3079 the TOSCA 1.0 specification, it is not required to put TOSCA definitions files into a special “Definitions”
3080 directory, but definitions YAML files can be placed into any directory within the CSAR file.

3081 6.2 TOSCA Meta File

3082 The **TOSCA.meta** file structure follows the exact same syntax as defined in the TOSCA 1.0 specification.
3083 However, it is only required to include `block_0` (see section 16.2 in [**TOSCA-1.0**]) with the **Entry-**
3084 **Definitions** keyword pointing to a valid TOSCA definitions YAML file that a TOSCA orchestrator should
3085 use as entry for parsing the contents of the overall CSAR file.

3086 Note that it is not required to explicitly list TOSCA definitions files in subsequent blocks of the
3087 **TOSCA.meta** file, but any TOSCA definitions files besides the one denoted by the **Entry-Definitions**
3088 keyword can be found by a TOSCA orchestrator by processing respective **imports** statements in the
3089 entry definitions file (or in recursively imported files).

3090 Note also that any additional artifact files (e.g. scripts, binaries, configuration files) do not have to be
3091 declared explicitly through blocks in the **TOSCA.meta** file. Instead, such artifacts will be fully described and
3092 pointed to by relative path names through artifact definitions in one of the TOSCA definitions files
3093 contained in the CSAR.

3094 Due to the simplified structure of the CSAR file and **TOSCA.meta** file compared to TOSCA 1.0, the **CSAR-**
3095 **Version** keyword listed in `block_0` of the meta-file is required to denote version **1.1**.

3096 6.2.1 Example

3097 The following listing represents a valid **TOSCA.meta** file according to this TOSCA Simple Profile
3098 specification.

```
TOSCA-Meta-File-Version: 1.0
CSAR-Version: 1.1
Created-By: OASIS TOSCA TC
Entry-Definitions: definitions/tosca_elk.yaml
```

3099

3100 This **TOSCA.meta** file indicates its simplified TOSCA Simple Profile structure by means of the **CSAR-**
3101 **Version** keyword with value **1.1**. The **Entry-Definitions** keyword points to a TOSCA definitions
3102 YAML file with the name **tosca_elk.yaml** which is contained in a directory called **definitions** within
3103 the root of the CSAR file.

3104 **6.3 Archive without TOSCA-Metadata**

3105 In case the archive doesn't contains a TOSCA-Metadata directory the archive is required to contains a
3106 single YAML file at the root of the archive (other templates may exists in sub-directories).

3107 This file must be a valid TOSCA definitions YAML file with the additional restriction that the metadata
3108 section (as defined in 3.9.3.2) is required and **template_name** and **template_version** metadata are also
3109 required.

3110 TOSCA processors should recognized this file as being the CSAR Entry-Definitions file. The CSAR-
3111 Version is defined by the **template_version** metadata section. The **Created-By** value is defined by the
3112 **template_author** metadata.

3113 **6.3.1 Example**

3114 The following represents a valid TOSCA template file acting as the CSAR Entry-Definitions file in an
3115 archive without TOSCA-Metadata directory.

```
tosca_definitions_version: tosca_simple_yaml_1_1

metadata:
  template_name: my_template
  template_author: OASIS TOSCA TC
  template_version: 1.0
```

3116

3117 7 TOSCA workflows

3118 TOSCA defines two different kinds of workflows that can be used to deploy (instantiate and start),
3119 manage at runtime or undeploy (stop and delete) a TOSCA topology: declarative workflows and
3120 imperative workflows. Declarative workflows are automatically generated by the TOSCA orchestrator
3121 based on the nodes, relationships, and groups defined in the topology. Imperative workflows are manually
3122 specified by the author of the topology and allows the specification of any use-case that has not been
3123 planned in the definition of node and relationships types or for advanced use-case (including reuse of
3124 existing scripts and workflows).

3125

3126 Workflows can be triggered on deployment of a topology (deploy workflow) on undeployment (undeploy
3127 workflow) or during runtime, manually, or automatically based on policies defined for the topology.

3128

3129 **Note:** The TOSCA orchestrators will execute a single workflow at a time on a topology to guarantee that
3130 the defined workflow can be consistent and behave as expected.

3131 7.1 Normative workflows

3132 TOSCA defines several normative workflows that are used to operate a Topology. That is, reserved
3133 names of workflows that should be preserved by TOSCA orchestrators and that, if specified in the
3134 topology will override the workflow generated by the orchestrator :

- 3135 • **deploy:** is the workflow used to instantiate and perform the initial deployment of the topology.
- 3136 • **undeploy:** is the workflow used to remove all instances of a topology.

3137 7.1.1 Notes

3138 Future versions of the specification will describe the normative naming and declarative generation of
3139 additional workflows used to operate the topology at runtime.

- 3140 • **scaling workflows:** defined for every scalable nodes or based on scaling policies
- 3141 • **auto-healing workflows:** defined in order to restart nodes that may have failed

3142 7.2 Declarative workflows

3143 Declarative workflows are the result of the weaving of topology's node, relationships, and groups
3144 workflows.

3145 The weaving process generates the workflow of every single node in the topology, insert operations from
3146 the relationships and groups and finally add ordering consideration. The weaving process will also take
3147 care of the specific lifecycle of some nodes and the TOSCA orchestrator is responsible to trigger errors or
3148 warnings in case the weaving cannot be processed or lead to cycles for example.

3149 This section aims to describe and explain how a TOSCA orchestrator will generate a workflow based on
3150 the topology entities (nodes, relationships and groups).

3151 7.2.1 Notes

3152 This section details specific constraints and considerations that applies during the weaving process.

3153 7.2.1.1 Orchestrator provided nodes lifecycle and weaving

3154 When a node is abstract the orchestrator is responsible for providing a valid matching resources for the
3155 node in order to deploy the topology. This consideration is also valid for dangling requirements (as they
3156 represents a quick way to define an actual node).

3157 The lifecycle of such nodes is the responsibility of the orchestrator and they may not answer to the
3158 normative TOSCA lifecycle. Their workflow is considered as "delegate" and acts as a black-box between
3159 the initial and started state in the install workflow and the started to deleted states in the uninstall
3160 workflow.

3161 If a relationship to some of this node defines operations or lifecycle dependency constraint that relies on
3162 intermediate states, the weaving SHOULD fail and the orchestrator SHOULD raise an error.

3163 7.2.2 Relationship impacts on topology weaving

3164 This section explains how relationships impacts the workflow generation to enable the composition of
3165 complex topologies.

3166 7.2.2.1 `tosca.relationships.DependsOn`

3167 The depends on relationship is used to establish a dependency from a node to another. A source node
3168 that depends on a target node will be created only after the other entity has been started.

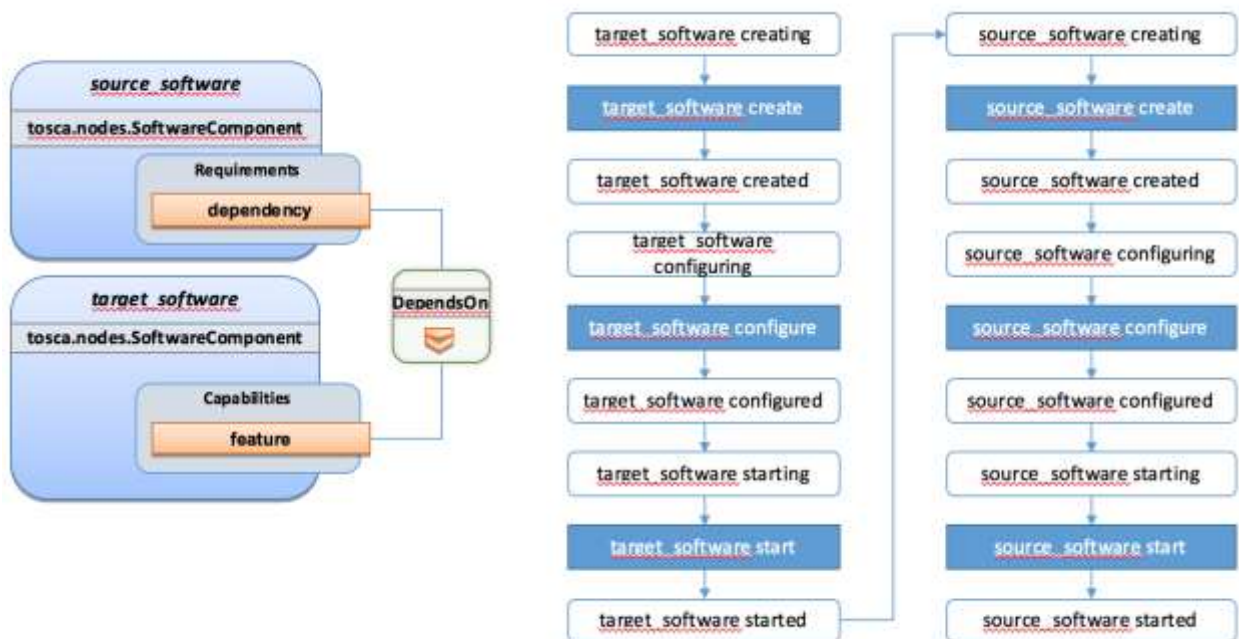
3169 7.2.2.2 Note

3170 DependsOn relationship SHOULD not be implemented. Even if the Configure interface can be
3171 implemented this is not considered as a best-practice. If you need specific implementation, please have a
3172 look at the ConnectsTo relationship.

3173 7.2.2.2.1 Example DependsOn

3174 This example show the usage of a generic DependsOn relationship between two custom software
3175 components.

3176

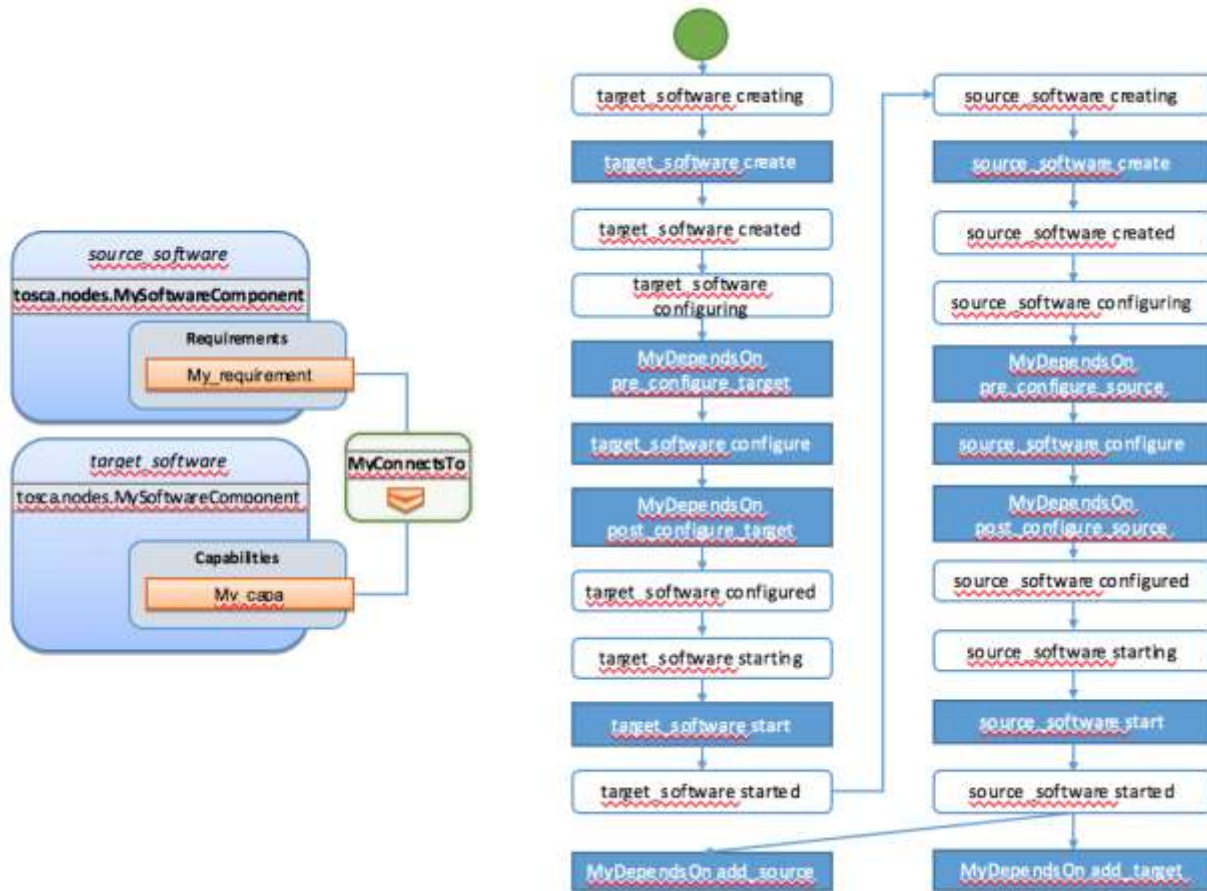


3177

3178 In this example the relationship configure interface doesn't define operations so they don't appear in the
3179 generated lifecycle.

3180 **7.2.2.3 tosca.relationships.ConnectsTo**

3181 The connects to relationship is similar to the DependsOn relationship except that it is intended to provide
3182 an implementation. The difference is more theoretical than practical but helps users to make an actual
3183 distinction from a meaning perspective.



3184

3185 **7.2.2.4 tosca.relationships.HostedOn**

3186 The hosted_on dependency relationship allows to define a hosting relationship between an entity and
3187 another. The hosting relationship has multiple impacts on the workflow and execution:

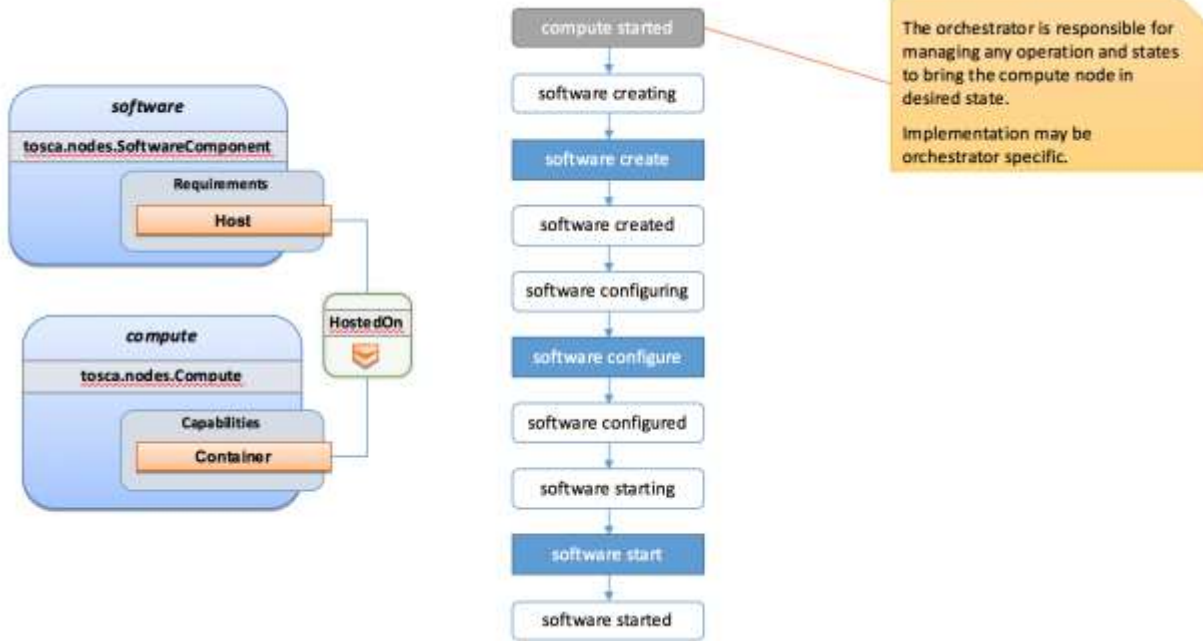
- 3188 • The implementation artifacts of the source node is executed on the same host as the one of the
3189 target node.
- 3190 • The create operation of the source node is executed only once the target node reach the started
3191 state.
- 3192 • When multiple nodes are hosted on the same host node, the defined operations will not be
3193 executed concurrently even if the theoretical workflow could allow it (actual generated workflow
3194 will avoid concurrency).

3195 **7.2.2.4.1 Example Software Component HostedOn Compute**

3196 This example explain the TOSCA weaving operation of a custom SoftwareComponent on a
3197 tosca.nodes.Compute instance. The compute node is an orchestrator provided node meaning that it's
3198 lifecycle is delegated to the orchestrator. This is a black-box and we just expect a started compute node
3199 to be provided by the orchestrator.

3200 The software node lifecycle operations will be executed on the Compute node (host) instance.

3201

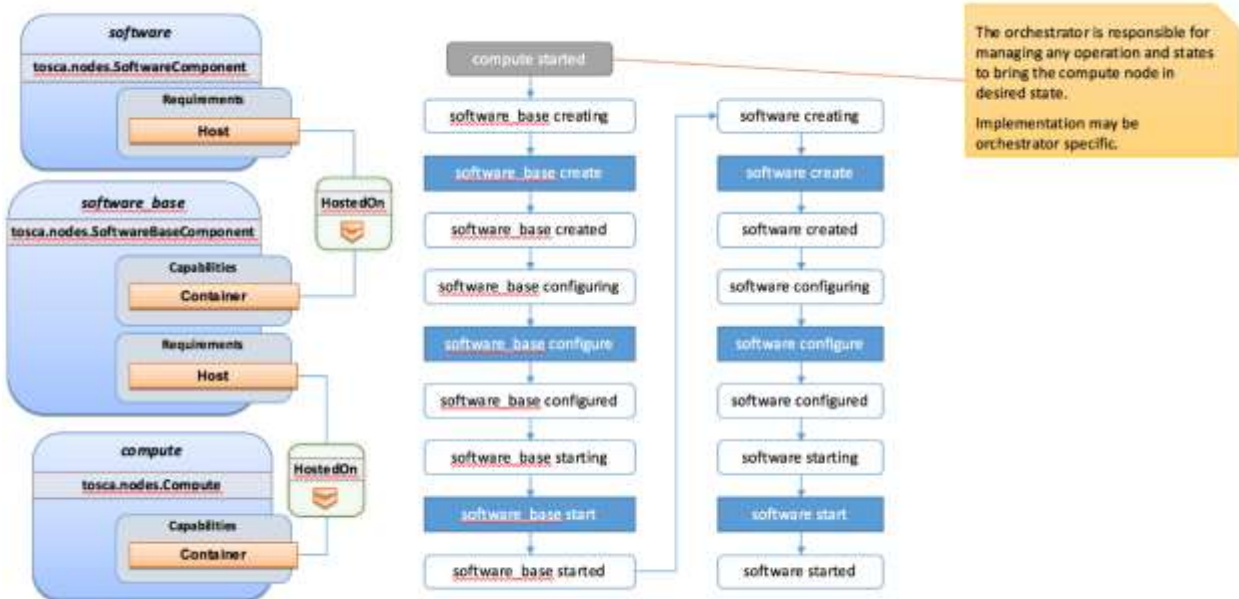


3202

3203 7.2.2.4.2 Example Software Component HostedOn Software Component

3204 Tosca allows some more complex hosting scenarios where a software component could be hosted on
3205 another software component.

3206



3207

3208 In such scenarios the software create operation is triggered only once the software_base node has
3209 reached the started state.

3210 **7.2.2.4.3 Example 2 Software Components HostedOn Compute**

3211 This example illustrate concurrency constraint introduced by the management of multiple nodes on a
3212 single compute.

3213 **7.2.3 Limitations**

3214 **7.2.3.1 Hosted nodes concurrency**

3215 TOSCA implementation currently does not allow concurrent executions of scripts implementation artifacts
3216 (shell, python, ansible, puppet, chef etc.) on a given host. This limitation is not applied on multiple hosts.
3217 This limitation is expressed through the HostedOn relationship limitation expressing that when multiple
3218 components are hosted on a given host node then their operations will not be performed concurrently
3219 (generated workflow will ensure that operations are not concurrent).

3220 **7.2.3.2 Dependent nodes concurrency**

3221 When a node depends on another node no operations will be processed concurrently. In some situations,
3222 especially when the two nodes lies on different hosts we could expect the create operation to be executed
3223 concurrently for performance optimization purpose. The current version of the specification will allow to
3224 use imperative workflows to solve this use-case. However, this scenario is one of the scenario that we
3225 want to improve and handle in the future through declarative workflows.

3226 **7.2.3.3 Target operations and get_attribute on source**

3227 The current ConnectsTo workflow implies that the target node is started before the source node is even
3228 created. This means that pre_configure_target and post_configure_target operations cannot use any
3229 input based on source attribute. It is however possible to refer to get_property inputs based on source
3230 properties. For advanced configurations the add_source operation should be used.

3231 Note also that future plans on declarative workflows improvements aims to solve this kind of issues while
3232 it is currently possible to use imperative workflows.

3233 **7.3 Imperative workflows**

3234 Imperative workflows are user defined and can define any really specific constraints and ordering of
3235 activities. They are really flexible and powerful and can be used for any complex use-case that cannot be
3236 solved in declarative workflows. However, they provide less reusability as they are defined for a specific
3237 topology rather than being dynamically generated based on the topology content.

3238 **7.3.1 Defining sequence of operations in an imperative workflow**

3239 Imperative workflow grammar defines two ways to define the sequence of operations in an imperative
3240 workflow:

- 3241 • Leverage the **on_success** definition to define the next steps that will be executed in parallel.
- 3242 • Leverage a sequence of activity in a step.

3243 **7.3.1.1 Using on_success to define steps ordering**

3244 The graph of workflow steps is build based on the values of **on_success** elements of the various defined
3245 steps. The graph is built based on the following rules:

- 3246 • All steps that defines an **on_success** operation must be executed before the next step can be
3247 executed. So if A and C defines an **on_success** operation to B, then B will be executed only
3248 when both A and C have been successfully executed.
- 3249 • The multiple nodes defined by an **on_success** construct can be executed in parallel.

- 3250 • Every step that doesn't have any predecessor is considered as an initial step and can run in parallel.
- 3251
- 3252 • Every step that doesn't define any successor is considered as final. When all the final nodes executions are completed then the workflow is considered as completed.
- 3253

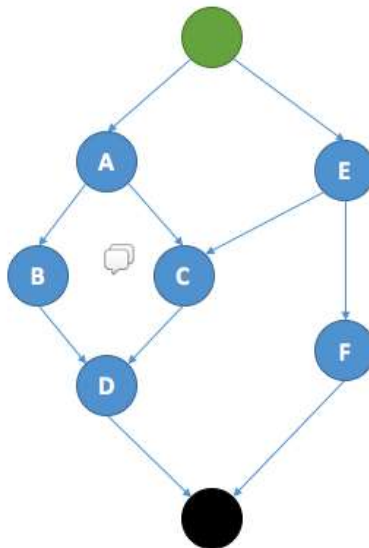
7.3.1.1.1 Example

The following example defines multiple steps and the **on_success** relationship between them.

```

topology_template:
  workflows:
    deploy:
      description: Workflow to deploy the application
      steps:
        A:
          on_success:
            - B
            - C
        B:
          on_success:
            - D
        C:
          on_success:
            - D
        D:
        E:
          on_success:
            - C
            - F
        F:
  
```

The following schema is the visualization of the above definition in term of sequencing of the steps.



3258

7.3.1.2 Define a sequence of activity on the same element

The step definition of a TOSCA imperative workflow allows multiple activities to be defined :

3261

```
workflows:
  my_workflow:
    steps:
      create_my_node:
        target: my_node
        activities:
          - set_state: creating
          - call_operation: tosca.interfaces.node.lifecycle.Standard.create
          - set_state: created
```

3262 The sequence defined here defines three different activities that will be performed in a sequential way.
3263 This is just equivalent to writing multiple steps chained by an on_success together :

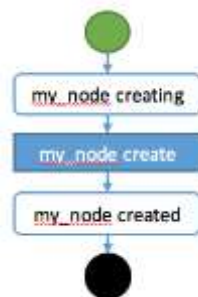
3264

3265

```
workflows:
  my_workflow:
    steps:
      creating_my_node:
        target: my_node
        activities:
          - set_state: creating
        on_success: create_my_node
      create_my_node:
        target: my_node
        activities:
          - call_operation: tosca.interfaces.node.lifecycle.Standard.create
        on_success: created_my_node
      created_my_node:
        target: my_node
        activities:
          - set_state: created
```

3266

3267 In both situations the resulting workflow is a sequence of activities:
3268



3269

3270 7.3.2 Definition of a simple workflow

3271 Imperative workflow allow user to define custom workflows allowing them to add operations that are not
3272 normative, or for example, to execute some operations in parallel when TOSCA would have performed
3273 sequential execution.

3274

3275 As Imperative workflows are related to a topology, adding a workflow is as simple as adding a workflows
3276 section to your topology template and specifying the workflow and the steps that compose it.

3277 7.3.2.1 Example: Adding a non-normative custom workflow

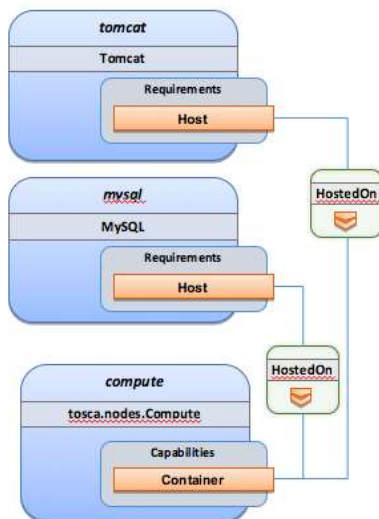
3278 This sample topology add a very simple custom workflow to trigger the mysql backup operation.

```
topology_template:  
  node_templates:  
    my_server:  
      type: toska.nodes.Compute  
    mysql:  
      type: toska.nodes.DBMS.MySQL  
      requirements:  
        - host: my_server  
      interfaces:  
        toska.interfaces.nodes.custom.Backup:  
          operations:  
            backup: backup.sh  
  workflows:  
    backup:  
      description: Performs a snapshot of the MySQL data.  
      steps:  
        my_step:  
          target: mysql  
          activities:  
            - call_operation: toska.interfaces.nodes.custom.Backup.backup
```

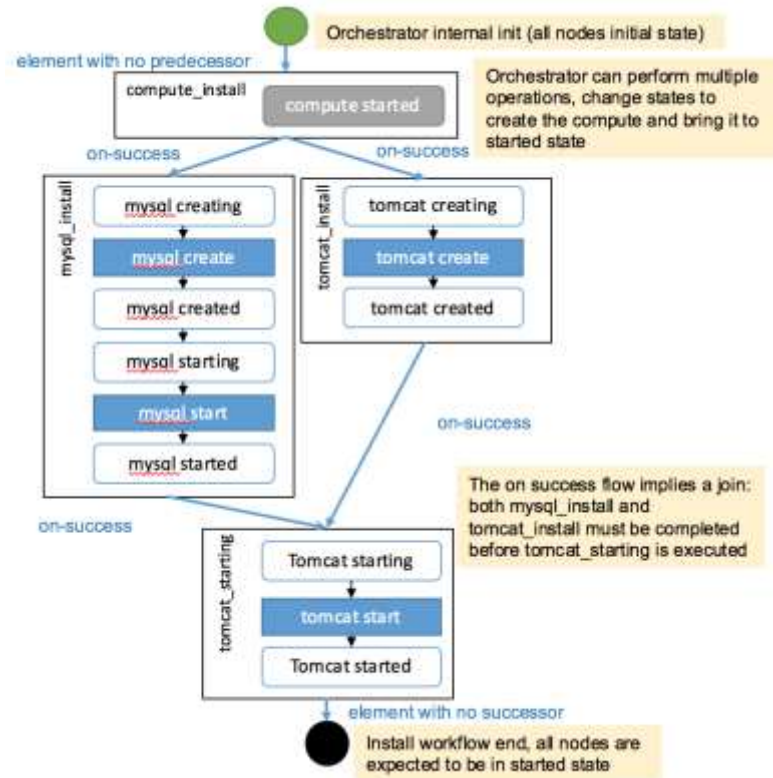
3279
3280 In such topology the TOSCA container will still use declarative workflow to generate the deploy and
3281 undeploy workflows as they are not specified and a backup workflow will be available for user to trigger.

3282 7.3.2.2 Example: Creating two nodes hosted on the same compute in parallel

3283 TOSCA declarative workflow generation constraint the workflow so that no operations are called in
3284 parallel on the same host. Looking at the following topology this means that the mysql and tomcat nodes
3285 will not be created in parallel but sequentially. This is fine in most of the situations as packet managers
3286 like apt or yum doesn't not support concurrency, however if both create operations performs a download
3287 of zip package from a server most of people will hope to do that in parallel in order to optimize throughput.



3289 Imperative workflows can help to solve this issue. Based on the above topology we will design a workflow
 3290 that will create tomcat and mysql in parallel but we will also ensure that tomcat is started after mysql
 3291 is started even if no relationship is defined between the components:
 3292



3293
 3294
 3295
 3296

To achieve such workflow, the following topology will be defined:

```

topology_template:
  node_templates:
    my_server:
      type: tosca.nodes.Compute
    mysql:
      type: tosca.nodes.DBMS.MySQL
      requirements:
        - host: my_server
    tomcat:
      type: tosca.nodes.WebServer.Tomcat
      requirements:
        - host: my_server
  workflows:
    deploy:
      description: Override the TOSCA declarative workflow with the following.
      steps:
        compute_install
        target: my_server
        activities:
          - delegate: deploy
  
```

```

    on_success:
      - mysql_install
      - tomcat_install
  tomcat_install:
    target: tomcat
    activities:
      - set_state: creating
      - call_operation: tosca.interfaces.node.lifecycle.Standard.create
      - set_state: created
    on_success:
      - tomcat_starting
  mysql_install:
    target: mysql
    activities:
      - set_state: creating
      - call_operation: tosca.interfaces.node.lifecycle.Standard.create
      - set_state: created
      - set_state: starting
      - call_operation: tosca.interfaces.node.lifecycle.Standard.start
      - set_state: started
    on_success:
      - tomcat_starting
  tomcat_starting:
    target: tomcat
    activities:
      - set_state: starting
      - call_operation: tosca.interfaces.node.lifecycle.Standard.start
      - set_state: started

```

3297

3298 **7.3.3 Specifying preconditions to a workflow**

3299 Pre conditions allows the TOSCA orchestrator to determine if a workflow can be executed based on the
 3300 states and attribute values of the topology's node. Preconditions must be added to the initial workflow.

3301 **7.3.3.1 Example : adding precondition to custom backup workflow**

3302 In this example we will use precondition so that we make sure that the mysql node is in the correct state
 3303 for a backup.

```

topology_template:
  node_templates:
    my_server:
      type: tosca.nodes.Compute
    mysql:
      type: tosca.nodes.DBMS.MySQL
      requirements:
        - host: my_server
      interfaces:
        tosca.interfaces.nodes.custom.Backup:
          operations:
            backup: backup.sh

```



```

workflows:
  backup:
    description: Performs a snapshot of the MySQL data.
    preconditions:
      - target: my_server
        condition:
          - assert:
              - state: [{equal: available}]
      - target: mysql
        condition:
          - assert:
              - state: [{valid_values: [started, available]}]
              - my_attribute: [{equal: ready }]
    steps:
      my_step:
        target: mysql
        activities:
          - call_operation: toska.interfaces.nodes.custom.Backup.backup

```

3304 When the backup workflow will be triggered (by user or policy) the TOSCA engine will first check that
 3305 preconditions are fulfilled. In this situation the engine will check that *my_server* node is in *available* state
 3306 AND that *mysql* node is in *started* OR *available* states AND that *mysql my_attribute* value is equal to
 3307 *ready*.

3308 7.3.4 Workflow reusability

3309 TOSCA allows the reusability of a workflow in other workflows. Such concepts can be achieved thanks to
 3310 the inline activity.

3311 7.3.4.1 Reusing a workflow to build multiple workflows

3312 The following example show how a workflow can inline an existing workflow and reuse it.

3313

```

topology_template:
  node_templates:
    my_server:
      type: toska.nodes.Compute
    mysql:
      type: toska.nodes.DBMS.MySQL
      requirements:
        - host: my_server
      interfaces:
        toska.interfaces.nodes.custom.Backup:
          operations:
            backup: backup.sh
  workflows:
    start_mysql:
      steps:
        start_mysql:
          target: mysql
          activities :
            - set_state: starting
            - call_operation: toska.interfaces.node.lifecycle.Standard.start
            - set_state: started
    stop_mysql:

```

```

steps:
  stop_mysql:
    target: mysql
    activities:
      - set_state: stopping
      - call_operation: toska.interfaces.node.lifecycle.Standard.stop
      - set_state: stopped

backup:
  description: Performs a snapshot of the MySQL data.
  preconditions:
    - target: my_server
      condition:
        - assert:
            - state: [{equal: available}]
    - target: mysql
      condition:
        - assert:
            - state: [{valid_values: [started, available]}]
            - my_attribute: [{equal: ready }]
  steps:
    backup_step:
      activities:
        - inline: stop
        - call_operation: toska.interfaces.nodes.custom.Backup.backup
        - inline: start

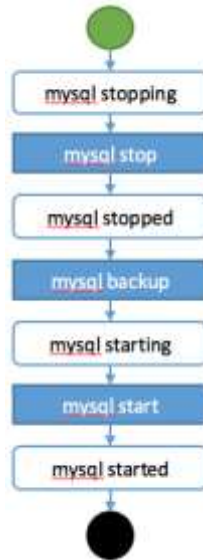
restart:
  steps:
    backup_step:
      activities:
        - inline: stop
        - inline: start

```

3314

3315 The example above defines three workflows and show how the start_mysql and stop_mysql workflows
 3316 are reused in the backup and restart workflows.

3317 Inlined workflows are inlined sequentially in the existing workflow for example the backup workflow would
 3318 look like this:



3319

3320 7.3.4.2 Inlining a complex workflow

3321 It is possible of course to inline more complex workflows. The following example defines an inlined
 3322 workflows with multiple steps including concurrent steps:

3323

```

topology_template:
  workflows:
    inlined_wf:
      steps:
        A:
          target: node_a
          activities:
            - call_operation: a
          on_success:
            - B
            - C
        B:
          target: node_a
          activities:
            - call_operation: b
          on_success:
            - D
        C:
          target: node_a
          activities:
            - call_operation: c
          on_success:
            - D
        D:
          target: node_a
          activities:
            - call_operation: d
        E:
          target: node_a
          activities:
            - call_operation: e
  
```

```

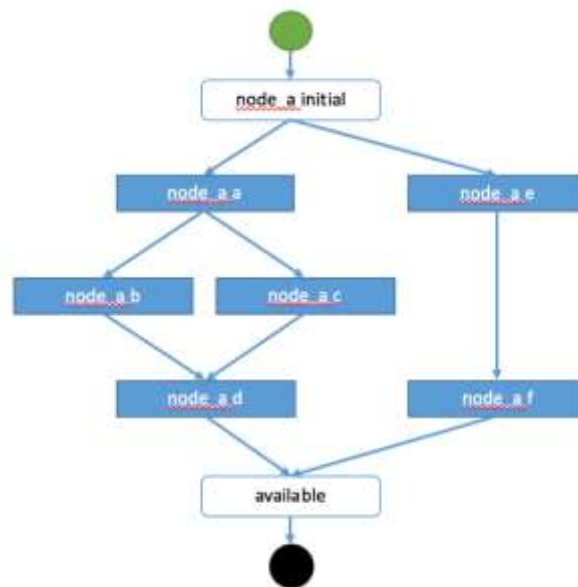
    on_success:
      - C
      - F
  F:
    target: node_a
    activities:
      - call_operation: f
main_workflow:
  steps:
    G:
      target: node_a
      activities:
        - set_state: initial
        - inline: inlined_wf
        - set_state: available

```

3324

3325 To describe the following workflow:

3326



3327

3328 7.3.5 Defining conditional logic on some part of the workflow

3329 Preconditions are used to validate if the workflow should be executed only for the initial workflow. If a
 3330 workflow that is inlined defines some preconditions these preconditions will be used at the instance level
 3331 to define if the operations should be executed or not on the defined instance.

3332

3333 This construct can be used to filter some steps on a specific instance or under some specific
 3334 circumstances or topology state.

3335

```

topology_template:
  node_templates:

```

```

my_server:
  type: tosca.nodes.Compute
cluster:
  type: tosca.nodes.DBMS.Cluster
  requirements:
    - host: my_server
  interfaces:
    tosca.interfaces.nodes.custom.Backup:
      operations:
        backup: backup.sh
workflows:
  backup:
    description: Performs a snapshot of the MySQL data.
    preconditions:
      - target: my_server
        condition:
          - assert:
              - state: [{equal: available}]
      - target: mysql
        condition:
          - assert:
              - state: [{valid_values: [started, available]}]
              - my_attribute: [{equal: ready }]
    steps:
      backup_step:
        target: cluster
        filter: # filter is a list of clauses. Matching between clauses is and.
          - or: # only one of sub-clauses must be true.
              - assert:
                  - foo: [{equals: true}]
              - assert:
                  - bar: [{greater_than: 2}, {less_than: 20}]
        activities:
          - call_operation: tosca.interfaces.nodes.custom.Backup.backup

```

3336

3337 **7.3.6 Define inputs for a workflow**

3338 Inputs can be defined in a workflow and will be provided in the execution context of the workflow. If an
 3339 operation defines a get_input function on one of its parameter the input will be retrieved from the workflow
 3340 input, and if not found from the topology inputs.

3341

3342 Workflow inputs will never be configured from policy triggered workflows and SHOULD be used only for
 3343 user triggered workflows. Of course operations can still refer to topology inputs or template properties or
 3344 attributes even in the context of a policy triggered workflow.

3345 **7.3.6.1 Example**

```

topology_template:
  node_templates:
    my_server:
      type: tosca.nodes.Compute
    mysql:
      type: tosca.nodes.DBMS.MySQL
      requirements:

```

```

- host: my_server
interfaces:
  toska.interfaces.nodes.custom.Backup:
    operations:
      backup:
        implementation: backup.sh
        inputs:
          storage_url: { get_input: storage_url }
workflows:
  backup:
    description: Performs a snapshot of the MySQL data.
    preconditions:
      - target: my_server
        valid_states: [available]
      - target: mysql
        valid_states: [started, available]
    attributes:
      my_attribute: [ready]
    inputs:
      storage_url:
        type: string
    steps:
      my_step:
        target: mysql
        activities:
          - call_operation: toska.interfaces.nodes.custom.Backup.backup

```

3346

3347 To trigger such a workflow, the TOSCA engine must allow user to provide inputs that match the given
3348 definitions.

3349 **7.3.7 Handle operation failure**

3350 By default, failure of any activity of the workflow will result in the failure of the workflow and will results in
3351 stopping the steps to be executed.

3352

3353 Exception: uninstall workflow operation failure SHOULD not prevent the other operations of the workflow
3354 to run (a failure in an uninstall script SHOULD not prevent from releasing resources from the cloud).

3355

3356 For any workflow other than install and uninstall failures may leave the topology in an unknown state. In
3357 such situation the TOSCA engine may not be able to orchestrate the deployment. Implementation of
3358 **on_failure** construct allows to execute rollback operations and reset the state of the affected entities
3359 back to an orchestrator known state.

3360 **7.3.7.1 Example**

```

topology_template:
  node_templates:
    my_server:
      type: toska.nodes.Compute
    mysql:
      type: toska.nodes.DBMS.MySQL
      requirements:
        - host: my_server

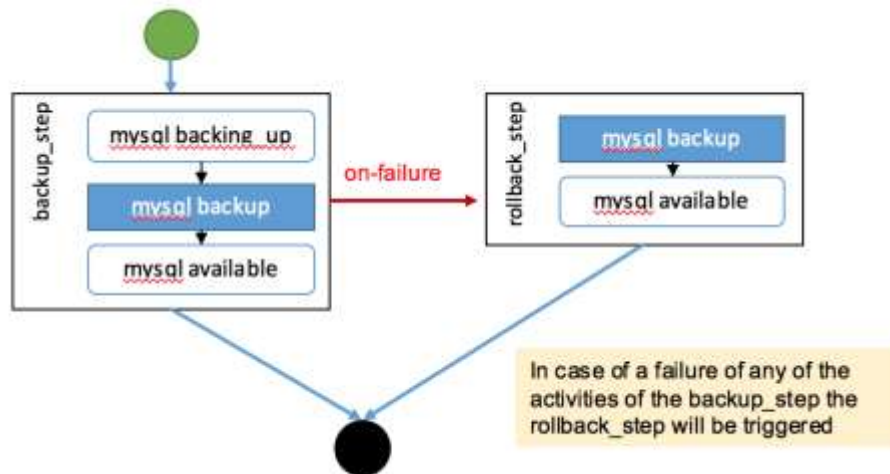
```

```

interfaces:
  tosca.interfaces.nodes.custom.Backup:
    operations:
      backup:
        implementation: backup.sh
        inputs:
          storage_url: { get_input: storage_url }
workflows:
  backup:
    steps:
      backup_step:
        target: mysql
        activities:
          - set_state: backing_up # this state is not a TOSCA known state
          - call_operation: tosca.interfaces.nodes.custom.Backup.backup
          - set_state: available # this state is known by TOSCA orchestrator
        on_failure:
          - rollback_step
      rollback_step:
        target: mysql
        activities:
          - call_operation: tosca.interfaces.nodes.custom.Backup.backup
          - set_state: available # this state is known by TOSCA orchestrator

```

3361



3362

3363

3364 7.4 Making declarative more flexible and imperative more generic

3365 TOSCA simple profile 1.1 version provides the genericity and reusability of declarative workflows that is
 3366 designed to address most of use-cases and the flexibility of imperative workflows to address more
 3367 complex or specific use-cases.

3368

3369 Each approach has some pros and cons and we are working so that the next versions of the specification
 3370 can improve the workflow usages to try to allow more flexibility in a more generic way. Two non-exclusive
 3371 leads are currently being discussed within the working group and may be included in the future versions
 3372 of the specification.

- 3373 • Improvement of the declarative workflows in order to allow people to extend the weaving logic of
 3374 TOSCA to fit some specific need.

- 3375 • Improvement of the imperative workflows in order to allow partial imperative workflows to be
3376 automatically included in declarative workflows based on specific constraints on the topology
3377 elements.

3378 Implementation of the improvements will be done by adding some elements to the specification and will
3379 not break compatibility with the current specification.

3380 **7.4.1.1 Notes**

- 3381 • The weaving improvement section is a Work in Progress and is not final in 1.1 version. The
3382 elements in this section are incomplete and may be subject to change in next specification
3383 version.
- 3384 • Moreover, the weaving improvements is one of the track of improvements. As describe improving
3385 the reusability of imperative workflow is another track (that may both co-exists in next
3386 specifications).

3387 **7.4.2 Weaving improvements**

3388 Making declarative better experimental option.

3389 **7.4.2.1 Node lifecycle definition**

3390 Node workflow is defined at the node type level. The node workflow definition is used to generate the
3391 declarative workflow of a given node.

3392 The `tosca.nodes.Root` type defines workflow steps for both the install workflow (used to instantiate or
3393 deploy a topology) and the uninstall workflow (used to destroy or undeploy a topology). The workflow is
3394 defined as follows:

3395

```
node_types:
  tosca.nodes.Root:
    workflows:
      install:
        steps:
          install_sequence:
            activities:
              - set_state: creating
              - call_operation: tosca.interfaces.node.lifecycle.Standard.create
              - set_state: created
              - set_state: configuring
              - call_operation: tosca.interfaces.node.lifecycle.Standard.configure
              - set_state: configured
              - set_state: starting
              - call_operation: tosca.interfaces.node.lifecycle.Standard.start
              - set_state: started
          uninstall:
            steps:
              uninstall_sequence:
                activities:
                  - set_state: stopping
                  - call_operation: tosca.interfaces.node.lifecycle.Standard.stop
                  - set_state: stopped
                  - set_state: deleting
                  - call_operation: tosca.interfaces.node.lifecycle.Standard.delete
                  - set_state: deleted
```


3396

3397 **7.4.2.2 Relationship lifecycle and weaving**

3398 While the workflow of a single node is quite simple the TOSCA weaving process is the real key element of
3399 declarative workflows. The process of weaving consist of the ability to create complex management
3400 workflows including dependency management in execution order between node operations, injection of
3401 operations to process specific instruction related to the connection to other nodes based the relationships
3402 and groups defined in a topology.

3403

3404 This section describes the relationship weaving and how the description at a template level can be
3405 translated on an instance level.

```
relationship_types:
  toska.relationships.ConnectsTo:
    workflow:
      install: # name of the workflow for wich the weaving has to be taken in
account
      source_weaving: # Instruct how to weave some tasks on the source workflow
(executed on SOURCE instance)
        - after: configuring # instruct that this operation should be weaved
after the target reach configuring state
          wait_target: created # add a join from a state of the target
          activity: toska.interfaces.relationships.Configure.pre_configure_source
        - before: configured # instruct that this operation should be weaved
before the target reach configured state
          activity:
toska.interfaces.relationships.Configure.post_configure_source
        - before: starting
          wait_target: started # add a join from a state of the target
        - after: started
          activity: toska.interfaces.relationships.Configure.add_target
      target_weaving: # Instruct how to weave some tasks on the target workflow
(executed on TARGET instance)
        - after: configuring # instruct that this operation should be weaved
after the target reach configuring state
          after_source: created # add a join from a state of the source
          activity: toska.interfaces.relationships.Configure.pre_configure_target
        - before: configured # instruct that this operation should be weaved
before the target reach configured state
          activity:
toska.interfaces.relationships.Configure.post_configure_target
        - after: started
          activity: toska.interfaces.relationships.Configure.add_source
```

3406

3407

8 TOSCA networking

3408
3409

Except for the examples, this section is **normative** and describes how to express and control the application centric network semantics available in TOSCA.

3410

8.1 Networking and Service Template Portability

3411
3412
3413
3414
3415

TOSCA Service Templates are application centric in the sense that they focus on describing application components in terms of their requirements and interrelationships. In order to provide cloud portability, it is important that a TOSCA Service Template avoid cloud specific requirements and details. However, at the same time, TOSCA must provide the expressiveness to control the mapping of software component connectivity to the network constructs of the hosting cloud.

3416

TOSCA Networking takes the following approach.

3417
3418
3419
3420
3421
3422
3423
3424
3425
3426
3427
3428
3429
3430
3431
3432
3433
3434
3435
3436
3437
3438
3439
3440

1. The application component connectivity semantics and expressed in terms of Requirements and Capabilities and the relationships between these. Service Template authors are able to express the interconnectivity requirements of their software components in an abstract, declarative, and thus highly portable manner.
2. The information provided in TOSCA is complete enough for a TOSCA implementation to fulfill the application component network requirements declaratively (i.e., it contains information such as communication initiation and layer 4 port specifications) so that the required network semantics can be realized on arbitrary network infrastructures.
3. TOSCA Networking provides full control of the mapping of software component interconnectivity to the networking constructs of the hosting cloud network independently of the Service Template, providing the required separation between application and network semantics to preserve Service Template portability.
4. Service Template authors have the choice of specifying application component networking requirements in the Service Template or completely separating the application component to network mapping into a separate document. This allows application components with explicit network requirements to express them while allowing users to control the complete mapping for all software components which may not have specific requirements. Usage of these two approaches is possible simultaneously and required to avoid having to re-write components network semantics as arbitrary sets of components are assembled into Service Templates.
5. Defining a set of network semantics which are expressive enough to address the most common application connectivity requirements while avoiding dependencies on specific network technologies and constructs. Service Template authors and cloud providers are able to express unique/non-portable semantics by defining their own specialized network Requirements and Capabilities.

3441

8.2 Connectivity Semantics

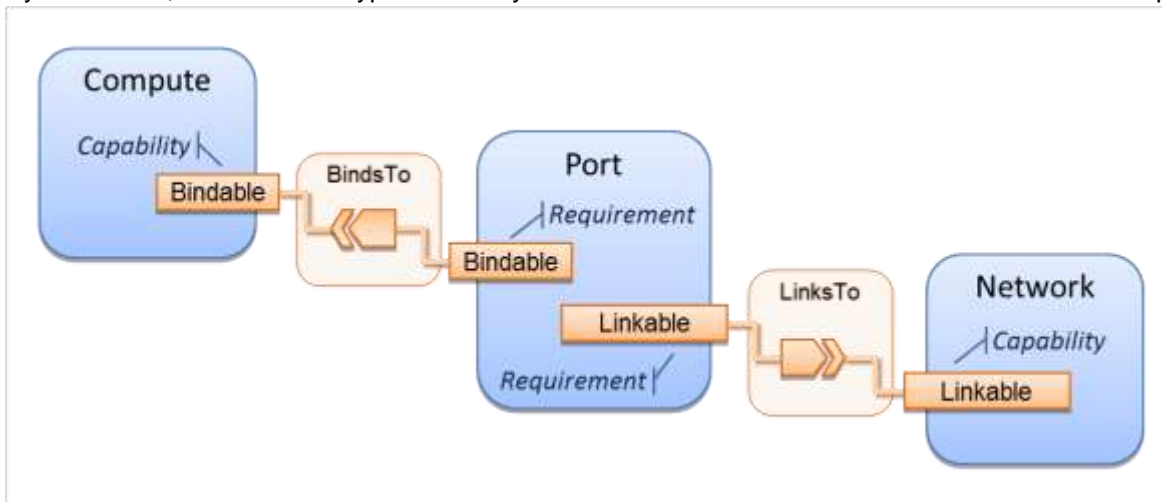
3442
3443
3444
3445
3446
3447
3448
3449

TOSCA's application centric approach includes the modeling of network connectivity semantics from an application component connectivity perspective. The basic premise is that applications contain components which need to communicate with other components using one or more endpoints over a network stack such as TCP/IP, where connectivity between two components is expressed as a <source component, source address, source port, target component, target address, target port> tuple. Note that source and target components are added to the traditional 4 tuple to provide the application centric information, mapping the network to the source or target component involved in the connectivity.

3450 Software components are expressed as Node Types in TOSCA which can express virtually any kind of
3451 concept in a TOSCA model. Node Types offering network based functions can model their connectivity
3452 using a special Endpoint Capability, [tosca.capabilities.Endpoint](#), designed for this purpose. Node Types
3453 which require an Endpoint can specify this as a TOSCA requirement. A special Relationship Type,
3454 [tosca.relationships.ConnectsTo](#), is used to implicitly or explicitly relate the source Node Type's endpoint
3455 to the required endpoint in the target node type. Since [tosca.capabilities.Endpoint](#) and
3456 [tosca.relationships.ConnectsTo](#) are TOSCA types, they can be used in templates and extended by
3457 subclassing in the usual ways, thus allowing the expression of additional semantics as needed.

3458 The following diagram shows how the TOSCA node, capability and relationship types enable modeling
3459 the application layer decoupled from the network model intersecting at the Compute node using the
3460 [Bindable](#) capability type.

3461 As you can see, the Port node type effectively acts a broker node between the Network node description



3462 and a host Compute node of an application.

3463 8.3 Expressing connectivity semantics

3464 This section describes how TOSCA supports the typical client/server and group communication
3465 semantics found in application architectures.

3466 8.3.1 Connection initiation semantics

3467 The [tosca.relationships.ConnectsTo](#) expresses that requirement that a source application component
3468 needs to be able to communicate with a target software component to consume the services of the target.
3469 [ConnectTo](#) is a component interdependency semantic in the most general sense and does not try imply
3470 how the communication between the source and target components is physically realized.

3471
3472 Application component intercommunication typically has conventions regarding which component(s)
3473 initiate the communication. Connection initiation semantics are specified in [tosca.capabilities.Endpoint](#).
3474 Endpoints at each end of the [tosca.relationships.ConnectsTo](#) must indicate identical connection initiation
3475 semantics.

3476
3477 The following sections describe the normative connection initiation semantics for the
3478 [tosca.relationships.ConnectsTo](#) Relationship Type.

3479 8.3.1.1 Source to Target

3480 The Source to Target communication initiation semantic is the most common case where the source
3481 component initiates communication with the target component in order to fulfill an instance of the

3482 `tosca.relationships.ConnectsTo` relationship. The typical case is a “client” component connecting to a
3483 “server” component where the client initiates a stream oriented connection to a pre-defined transport
3484 specific port or set of ports.

3485
3486 It is the responsibility of the TOSCA implementation to ensure the source component has a suitable
3487 network path to the target component and that the ports specified in the respective
3488 `tosca.capabilities.Endpoint` are not blocked. The TOSCA implementation may only represent state of the
3489 `tosca.relationships.ConnectsTo` relationship as fulfilled after the actual network communication is enabled
3490 and the source and target components are in their operational states.

3491
3492 Note that the connection initiation semantic only impacts the fulfillment of the actual connectivity and does
3493 not impact the node traversal order implied by the `tosca.relationships.ConnectsTo` Relationship Type.

3494 **8.3.1.2 Target to Source**

3495 The Target to Source communication initiation semantic is a less common case where the target
3496 component initiates communication with the source component in order to fulfill an instance of the
3497 `tosca.relationships.ConnectsTo` relationship. This “reverse” connection initiation direction is typically
3498 required due to some technical requirements of the components or protocols involved, such as the
3499 requirement that SSH must only be initiated from target component in order to fulfill the services required
3500 by the source component.

3501
3502 It is the responsibility of the TOSCA implementation to ensure the source component has a suitable
3503 network path to the target component and that the ports specified in the respective
3504 `tosca.capabilities.Endpoint` are not blocked. The TOSCA implementation may only represent state of the
3505 `tosca.relationships.ConnectsTo` relationship as fulfilled after the actual network communication is enabled
3506 and the source and target components are in their operational states.

3507
3508 Note that the connection initiation semantic only impacts the fulfillment of the actual connectivity and does
3509 not impact the node traversal order implied by the `tosca.relationships.ConnectsTo` Relationship Type.

3510 **8.3.1.3 Peer-to-Peer**

3511 The Peer-to-Peer communication initiation semantic allows any member of a group to initiate
3512 communication with any other member of the same group at any time. This semantic typically appears in
3513 clustering and distributed services where there is redundancy of components or services.

3514
3515 It is the responsibility of the TOSCA implementation to ensure the source component has a suitable
3516 network path between all the member component instances and that the ports specified in the respective
3517 `tosca.capabilities.Endpoint` are not blocked, and the appropriate multicast communication, if necessary,
3518 enabled. The TOSCA implementation may only represent state of the `tosca.relationships.ConnectsTo`
3519 relationship as fulfilled after the actual network communication is enabled such that at least one-member
3520 component of the group may reach any other member component of the group.

3521
3522 Endpoints specifying the Peer-to-Peer initiation semantic need not be related with a
3523 `tosca.relationships.ConnectsTo` relationship for the common case where the same set of component
3524 instances must communicate with each other.

3525
3526 Note that the connection initiation semantic only impacts the fulfillment of the actual connectivity and does
3527 not impact the node traversal order implied by the `tosca.relationships.ConnectsTo` Relationship Type.

3528 **8.3.2 Specifying layer 4 ports**

3529 TOSCA Service Templates must express enough details about application component
3530 intercommunication to enable TOSCA implementations to fulfill these communication semantics in the
3531 network infrastructure. TOSCA currently focuses on TCP/IP as this is the most pervasive in today's cloud
3532 infrastructures. The layer 4 ports required for application component intercommunication are specified in
3533 `tosca.capabilities.Endpoint`. The union of the port specifications of both the source and target
3534 `tosca.capabilities.Endpoint` which are part of the `tosca.relationships.ConnectsTo` Relationship Template
3535 are interpreted as the effective set of ports which must be allowed in the network communication.

3536

3537 The meaning of Source and Target port(s) corresponds to the direction of the respective
3538 `tosca.relationships.ConnectsTo`.

3539 **8.4 Network provisioning**

3540 **8.4.1 Declarative network provisioning**

3541 TOSCA orchestrators are responsible for the provisioning of the network connectivity for declarative
3542 TOSCA Service Templates (Declarative TOSCA Service Templates don't contain explicit plans). This
3543 means that the TOSCA orchestrator must be able to infer a suitable logical connectivity model from the
3544 Service Template and then decide how to provision the logical connectivity, referred to as "fulfillment", on
3545 the available underlying infrastructure. In order to enable fulfillment, sufficient technical details still must
3546 be specified, such as the required protocols, ports and QOS information. TOSCA connectivity types, such
3547 as `tosca.capabilities.Endpoint`, provide well defined means to express these details.

3548 **8.4.2 Implicit network fulfillment**

3549 TOSCA Service Templates are by default network agnostic. TOSCA's application centric approach only
3550 requires that a TOSCA Service Template contain enough information for a TOSCA orchestrator to infer
3551 suitable network connectivity to meet the needs of the application components. Thus Service Template
3552 designers are not required to be aware of or provide specific requirements for underlying networks. This
3553 approach yields the most portable Service Templates, allowing them to be deployed into any
3554 infrastructure which can provide the necessary component interconnectivity.

3555 **8.4.3 Controlling network fulfillment**

3556 TOSCA provides mechanisms for providing control over network fulfillment.

3557 This mechanism allows the application network designer to express in service template or network
3558 template how the networks should be provisioned.

3559

3560 For the use cases described below let's assume we have a typical 3-tier application which is consisting of
3561 FE (frontend), BE (backend) and DB (database) tiers. The simple application topology diagram can be
3562 shown below:

3563

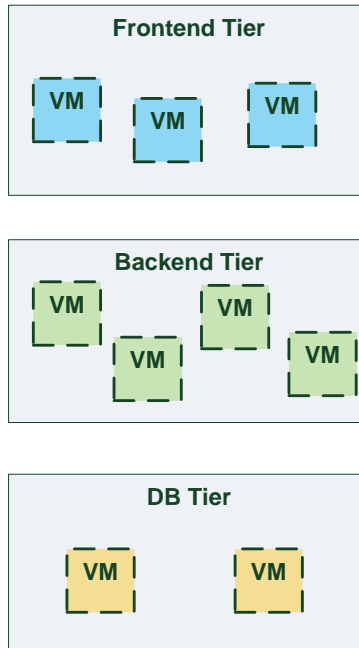


Figure-5: Typical 3-Tier Network

3564

3565

3566 8.4.3.1 Use case: OAM Network

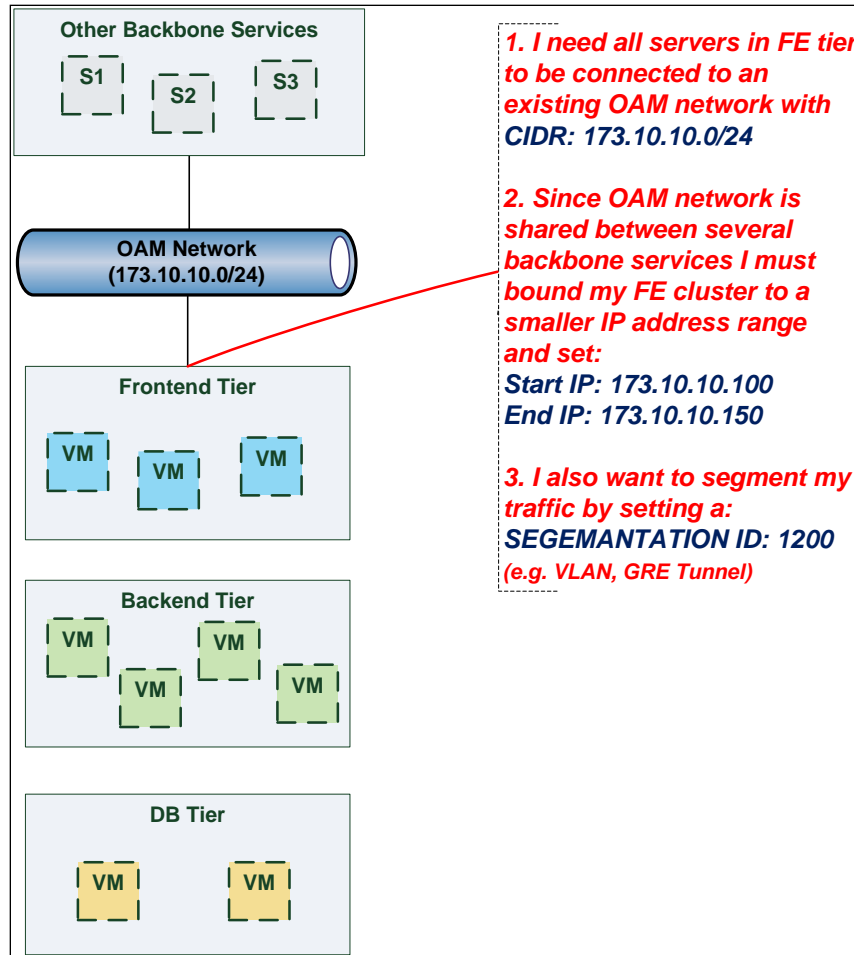
3567 When deploying an application in service provider's on-premise cloud, it's very common that one or more
 3568 of the application's services should be accessible from an ad-hoc OAM (Operations, Administration and
 3569 Management) network which exists in the service provider backbone.

3570

3571 As an application network designer, I'd like to express in my TOSCA network template (which
 3572 corresponds to my TOSCA service template) the network CIDR block, start ip, end ip and segmentation
 3573 ID (e.g. VLAN id).

3574 The diagram below depicts a typical 3-tiers application with specific networking requirements for its FE
 3575 tier server cluster:

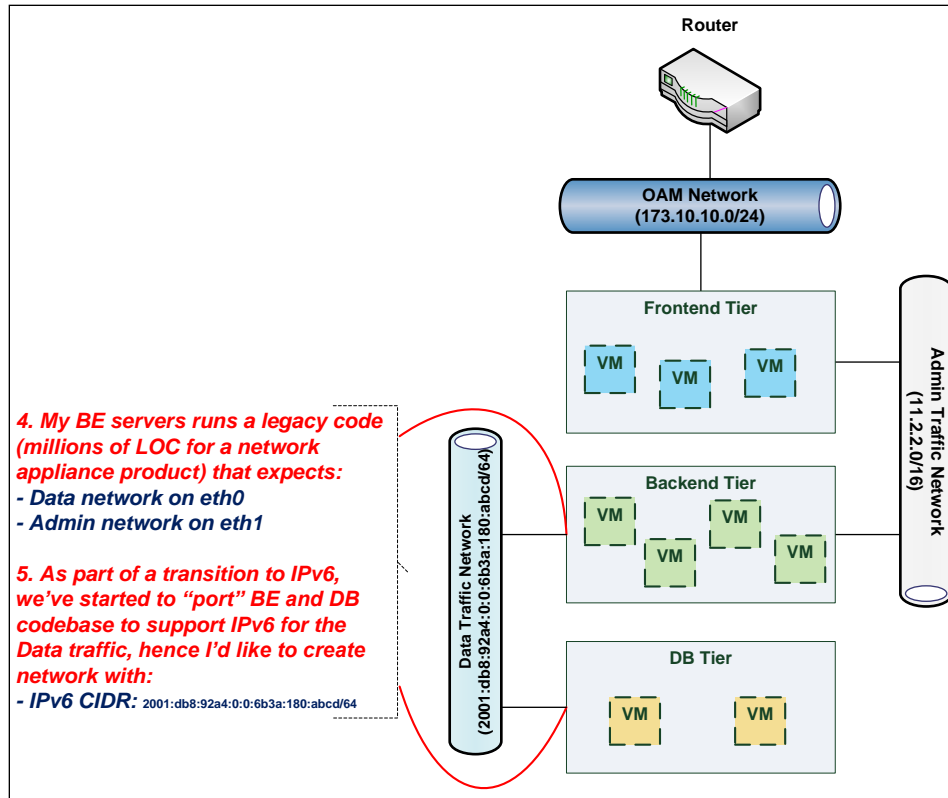
3576



3577

3578 **8.4.3.2 Use case: Data Traffic network**

3579 The diagram below defines a set of networking requirements for the backend and DB tiers of the 3-tier
3580 app mentioned above.



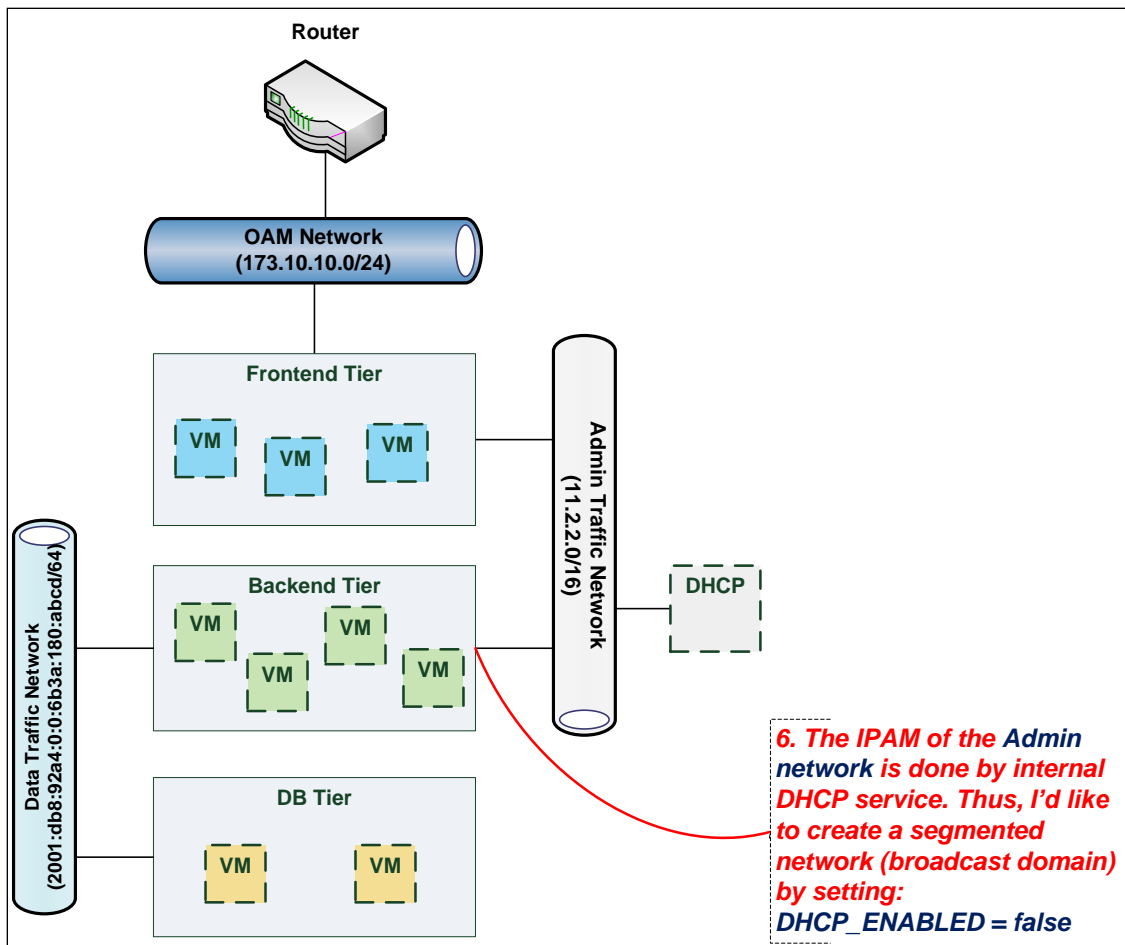
3581

3582 **8.4.3.3 Use case: Bring my own DHCP**

3583 The same 3-tier app requires for its admin traffic network to manage the IP allocation by its own DHCP
 3584 which runs autonomously as part of application domain.

3585

3586 For this purpose, the app network designer would like to express in TOSCA that the underlying
 3587 provisioned network will be set with DHCP_ENABLED=false. See this illustrated in the figure below:



3589 8.5 Network Types

3590 8.5.1 tosca.nodes.network.Network

3591 The TOSCA Network node represents a simple, logical network service.

Shorthand Name	Network
Type Qualified Name	tosca:Network
Type URI	tosca.nodes.network.Network

3592 8.5.1.1 Properties

Name	Required	Type	Constraints	Description
ip_version	no	integer	valid_values: [4, 6] default: 4	The IP version of the requested network
cidr	no	string	None	The cidr block of the requested network
start_ip	no	string	None	The IP address to be used as the 1 st one in a pool of addresses derived from the cidr block full IP range
end_ip	no	string	None	The IP address to be used as the last one in a pool of addresses derived from the cidr block full IP range

Name	Required	Type	Constraints	Description
gateway_ip	no	string	None	The gateway IP address.
network_name	no	string	None	An Identifier that represents an existing Network instance in the underlying cloud infrastructure – OR – be used as the name of the new created network. <ul style="list-style-type: none"> If network_name is provided along with network_id they will be used to uniquely identify an existing network and not creating a new one, means all other possible properties are not allowed. network_name should be more convenient for using. But in case that network name uniqueness is not guaranteed then one should provide a network_id as well.
network_id	no	string	None	An Identifier that represents an existing Network instance in the underlying cloud infrastructure. This property is mutually exclusive with all other properties except network_name. <ul style="list-style-type: none"> Appearance of network_id in network template instructs the Tosca container to use an existing network instead of creating a new one. network_name should be more convenient for using. But in case that network name uniqueness is not guaranteed then one should add a network_id as well. network_name and network_id can be still used together to achieve both uniqueness and convenient.
segmentation_id	no	string	None	A segmentation identifier in the underlying cloud infrastructure (e.g., VLAN id, GRE tunnel id). If the segmentation_id is specified, the network_type or physical_network properties should be provided as well.
network_type	no	string	None	Optionally, specifies the nature of the physical network in the underlying cloud infrastructure. Examples are flat, vlan, gre or vxlan. For flat and vlan types, physical_network should be provided too.
physical_network	no	string	None	Optionally, identifies the physical network on top of which the network is implemented, e.g. physnet1. This property is required if network_type is flat or vlan.
dhcp_enabled	no	boolean	default: true	Indicates the TOSCA container to create a virtual network instance with or without a DHCP service.

3593 8.5.1.2 Attributes

Name	Required	Type	Constraints	Description
segmentation_id	no	string	None	The actual <i>segmentation_id</i> that is been assigned to the network by the underlying cloud infrastructure.

3594 8.5.1.3 Definition

```
tosca.nodes.network.Network:
  derived_from: tosca.nodes.Root
  properties:
```

```

ip_version:
  type: integer
  required: false
  default: 4
  constraints:
    - valid_values: [ 4, 6 ]
cidr:
  type: string
  required: false
start_ip:
  type: string
  required: false
end_ip:
  type: string
  required: false
gateway_ip:
  type: string
  required: false
network_name:
  type: string
  required: false
network_id:
  type: string
  required: false
segmentation_id:
  type: string
  required: false
network_type:
  type: string
  required: false
physical_network:
  type: string
  required: false
capabilities:
  link:
    type: tosca.capabilities.network.Linkable

```

3595 **8.5.2 tosca.nodes.network.Port**

3596 The TOSCA **Port** node represents a logical entity that associates between Compute and Network
 3597 normative types.

3598 The Port node type effectively represents a single virtual NIC on the Compute node instance.

Shorthand Name	Port
Type Qualified Name	tosca:Port
Type URI	tosca.nodes.network.Port

3599 8.5.2.1 Properties

Name	Required	Type	Constraints	Description
ip_address	no	string	None	Allow the user to set a fixed IP address. Note that this address is a request to the provider which they will attempt to fulfill but may not be able to dependent on the network the port is associated with.
order	no	integer	greater_or_equal : 0 default: 0	The order of the NIC on the compute instance (e.g. eth2). Note: when binding more than one port to a single compute (aka multi vNICs) and ordering is desired, it is <i>*mandatory*</i> that all ports will be set with an order value and. The <i>order</i> values must represent a positive, arithmetic progression that starts with 0 (e.g. 0, 1, 2, ..., n).
is_default	no	boolean	default: false	Set is_default =true to apply a default gateway route on the running compute instance to the associated network gateway. Only one port that is associated to single compute node can set as default=true.
ip_range_start	no	string	None	Defines the starting IP of a range to be allocated for the compute instances that are associated by this Port. Without setting this property the IP allocation is done from the entire CIDR block of the network.
ip_range_end	no	string	None	Defines the ending IP of a range to be allocated for the compute instances that are associated by this Port. Without setting this property the IP allocation is done from the entire CIDR block of the network.

3600 8.5.2.2 Attributes

Name	Required	Type	Constraints	Description
ip_address	no	string	None	The IP address would be assigned to the associated compute instance.

3601 8.5.2.3 Definition

```
tosca.nodes.network.Port:
  derived_from: tosca.nodes.Root
  properties:
    ip_address:
      type: string
      required: false
    order:
      type: integer
```

```

    required: true
    default: 0
    constraints:
      - greater_or_equal: 0
  is_default:
    type: boolean
    required: false
    default: false
  ip_range_start:
    type: string
    required: false
  ip_range_end:
    type: string
    required: false
  requirements:
    - link:
        capability: tosca.capabilities.network.Linkable
        relationship: tosca.relationships.network.Linksto
    - binding:
        capability: tosca.capabilities.network.Bindable
        relationship: tosca.relationships.network.BindsTo

```

3602 8.5.3 [tosca.capabilities.network.Linkable](#)

3603 A node type that includes the Linkable capability indicates that it can be pointed to by a
 3604 [tosca.relationships.network.Linksto](#) relationship type.

Shorthand Name	Linkable
Type Qualified Name	tosca::Linkable
Type URI	tosca.capabilities.network.Linkable

3605 8.5.3.1 Properties

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

3606 8.5.3.2 Definition

```

tosca.capabilities.network.Linkable:
  derived_from: tosca.capabilities.Node

```

3607 8.5.4 [tosca.relationships.network.Linksto](#)

3608 This relationship type represents an association relationship between Port and Network node types.

Shorthand Name	Linksto
Type Qualified Name	tosca:Linksto
Type URI	tosca.relationships.network.Linksto

3609 8.5.4.1 Definition

```

tosca.relationships.network.Linksto:

```

```
derived_from: tosca.relationships.DependsOn
valid_target_types: [ tosca.capabilities.network.Linkable ]
```

3610 **8.5.5 [tosca.relationships.network.BindsTo](#)**

3611 This type represents a network association relationship between Port and Compute node types.

Shorthand Name	network.BindsTo
Type Qualified Name	tosca:BindsTo
Type URI	tosca.relationships.network.BindsTo

3612 **8.5.5.1 Definition**

```
tosca.relationships.network.BindsTo:
  derived_from: tosca.relationships.DependsOn
  valid_target_types: [ tosca.capabilities.network.Bindable ]
```

3613 **8.6 Network modeling approaches**

3614 **8.6.1 Option 1: Specifying a network outside the application's Service** 3615 **Template**

3616 This approach allows someone who understands the application's networking requirements, mapping the
3617 details of the underlying network to the appropriate node templates in the application.

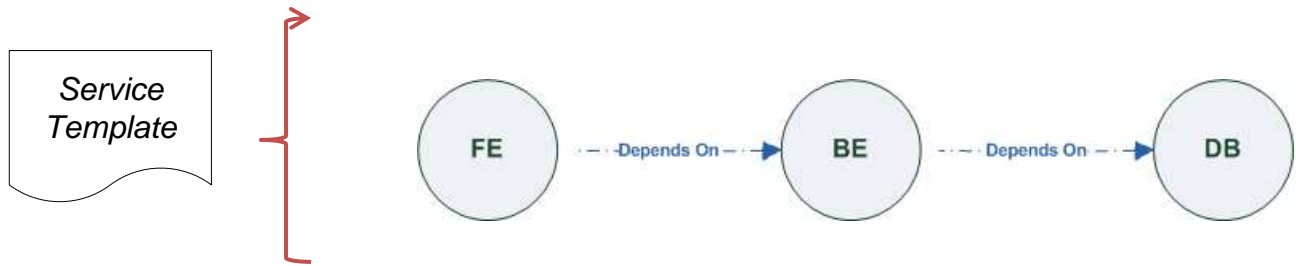
3618
3619 The motivation for this approach is providing the application network designer a fine-grained control on
3620 how networks are provisioned and stitched to its application by the TOSCA orchestrator and underlying
3621 cloud infrastructure while still preserving the portability of his service template. Preserving the portability
3622 means here not doing any modification in service template but just "plug-in" the desired network
3623 modeling. The network modeling can reside in the same service template file but the best practice should
3624 be placing it in a separated self-contained network template file.

3625
3626 This "pluggable" network template approach introduces a new normative node type called Port, capability
3627 called [tosca.capabilities.network.Linkable](#) and relationship type called
3628 [tosca.relationships.network.LinksTo](#).

3629 The idea of the Port is to elegantly associate the desired compute nodes with the desired network nodes
3630 while not "touching" the compute itself.

3631
3632 The following diagram series demonstrate the plug-ability strength of this approach.

3633 Let's assume an application designer has modeled a service template as shown in Figure 1 that
3634 describes the application topology nodes (compute, storage, software components, etc.) with their
3635 relationships. The designer ideally wants to preserve this service template and use it in any cloud
3636 provider environment without any change.

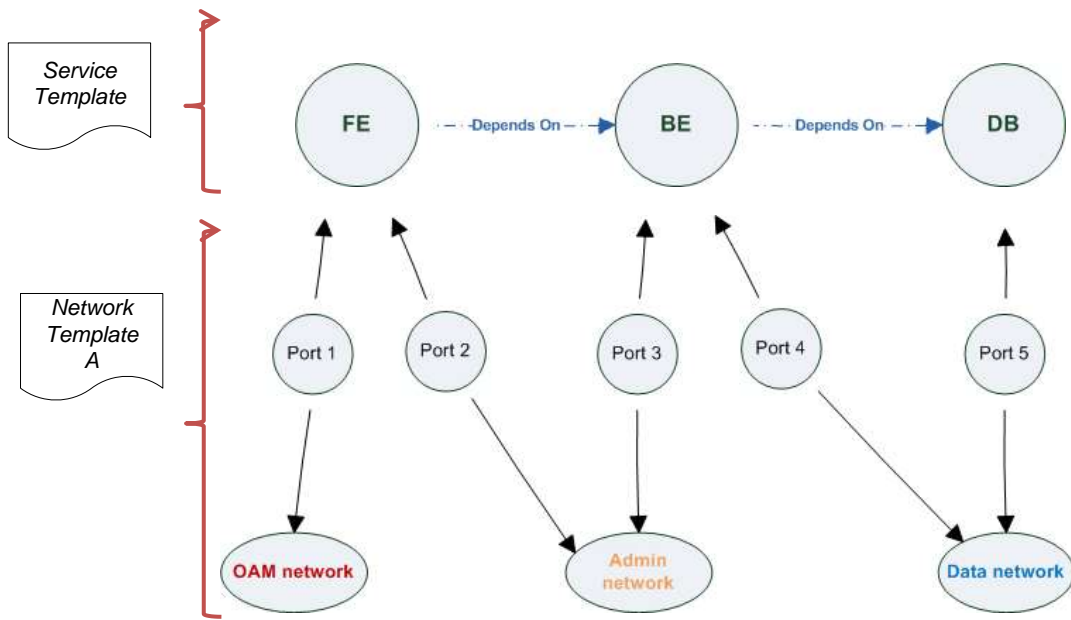


3637
3638

Figure-6: Generic Service Template

3639 When the application designer comes to consider its application networking requirement they typically call
3640 the network architect/designer from their company (who has the correct expertise).

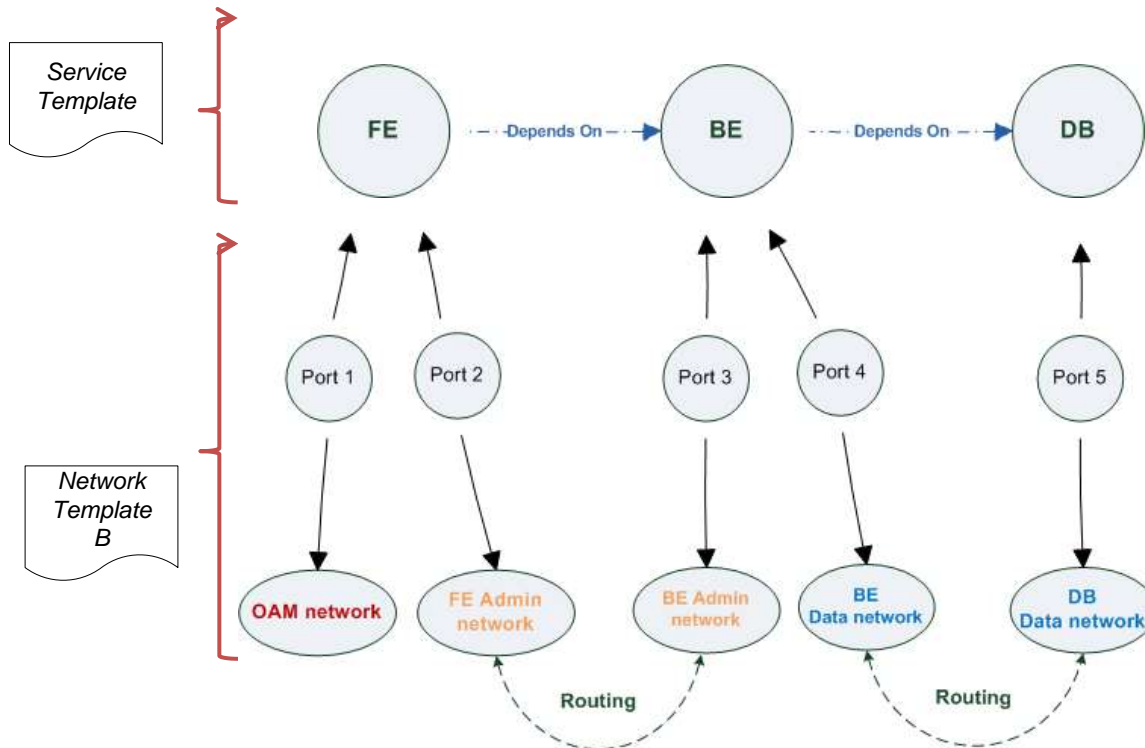
3641 The network designer, after understanding the application connectivity requirements and optionally the
3642 target cloud provider environment, is able to model the network template and plug it to the service
3643 template as shown in Figure 2:



3644
3645

Figure-7: Service template with network template A

3646 When there's a new target cloud environment to run the application on, the network designer is simply
3647 creates a new network template B that corresponds to the new environmental conditions and provide it to
3648 the application designer which packs it into the application CSAR.



3649

3650

Figure-8: Service template with network template B

3651 The node templates for these three networks would be defined as follows:

```

node_templates:
  frontend:
    type: toasca.nodes.Compute
    properties: # omitted for brevity

  backend:
    type: toasca.nodes.Compute
    properties: # omitted for brevity

  database:
    type: toasca.nodes.Compute
    properties: # omitted for brevity

  oam_network:
    type: toasca.nodes.network.Network
    properties: # omitted for brevity

  admin_network:
    type: toasca.nodes.network.Network
    properties: # omitted for brevity

  data_network:
    type: toasca.nodes.network.Network
    properties: # omitted for brevity

# ports definition
fe_oam_net_port:
  type: toasca.nodes.network.Port

```



```

properties:
  is_default: true
  ip_range_start: { get_input: fe_oam_net_ip_range_start }
  ip_range_end: { get_input: fe_oam_net_ip_range_end }
requirements:
  - link: oam_network
  - binding: frontend

fe_admin_net_port:
  type: tosca.nodes.network.Port
  requirements:
    - link: admin_network
    - binding: frontend

be_admin_net_port:
  type: tosca.nodes.network.Port
  properties:
    order: 0
  requirements:
    - link: admin_network
    - binding: backend

be_data_net_port:
  type: tosca.nodes.network.Port
  properties:
    order: 1
  requirements:
    - link: data_network
    - binding: backend

db_data_net_port:
  type: tosca.nodes.network.Port
  requirements:
    - link: data_network
    - binding: database

```

3652 **8.6.2 Option 2: Specifying network requirements within the application's**
3653 **Service Template**

3654 This approach allows the Service Template designer to map an endpoint to a logical network.

3655 The use case shown below examines a way to express in the TOSCA YAML service template a typical 3-
3656 tier application with their required networking modeling:

```

node_templates:
  frontend:
    type: tosca.nodes.Compute
    properties: # omitted for brevity
    requirements:
      - network_oam: oam_network
      - network_admin: admin_network
  backend:
    type: tosca.nodes.Compute
    properties: # omitted for brevity

```

```
requirements:
  - network_admin: admin_network
  - network_data: data_network

database:
  type: toska.nodes.Compute
  properties: # omitted for brevity
  requirements:
    - network_data: data_network

oam_network:
  type: toska.nodes.network.Network
  properties:
    ip_version: { get_input: oam_network_ip_version }
    cidr: { get_input: oam_network_cidr }
    start_ip: { get_input: oam_network_start_ip }
    end_ip: { get_input: oam_network_end_ip }

admin_network:
  type: toska.nodes.network.Network
  properties:
    ip_version: { get_input: admin_network_ip_version }
    dhcp_enabled: { get_input: admin_network_dhcp_enabled }

data_network:
  type: toska.nodes.network.Network
  properties:
    ip_version: { get_input: data_network_ip_version }
    cidr: { get_input: data_network_cidr }
```

3657

3658 9 Non-normative type definitions

3659 This section defines **non-normative** types which are used only in examples and use cases in this
3660 specification and are included only for completeness for the reader. Implementations of this specification
3661 are not required to support these types for conformance.

3662 9.1 Artifact Types

3663 This section contains are non-normative Artifact Types used in use cases and examples.

3664 9.1.1 `tosca.artifacts.Deployment.Image.Container.Docker`

3665 This artifact represents a Docker “image” (a TOSCA deployment artifact type) which is a binary comprised
3666 of one or more (a union of read-only and read-write) layers created from snapshots within the underlying
3667 Docker **Union File System**.

3668 9.1.1.1 Definition

```
tosca.artifacts.Deployment.Image.Container.Docker:  
  derived_from: tosca.artifacts.Deployment.Image  
  description: Docker Container Image
```

3669 9.1.2 `tosca.artifacts.Deployment.Image.VM.ISO`

3670 A Virtual Machine (VM) formatted as an ISO standard disk image.

3671 9.1.2.1 Definition

```
tosca.artifacts.Deployment.Image.VM.ISO:  
  derived_from: tosca.artifacts.Deployment.Image.VM  
  description: Virtual Machine (VM) image in ISO disk format  
  mime_type: application/octet-stream  
  file_ext: [ iso ]
```

3672 9.1.3 `tosca.artifacts.Deployment.Image.VM.QCOW2`

3673 A Virtual Machine (VM) formatted as a QEMU emulator version 2 standard disk image.

3674 9.1.3.1 Definition

```
tosca.artifacts.Deployment.Image.VM.QCOW2:  
  derived_from: tosca.artifacts.Deployment.Image.VM  
  description: Virtual Machine (VM) image in QCOW v2 standard disk format  
  mime_type: application/octet-stream  
  file_ext: [ qcow2 ]
```

3675 9.2 Capability Types

3676 This section contains are non-normative Capability Types used in use cases and examples.

3677 9.2.1 `tosca.capabilities.Container.Docker`

3678 The type indicates capabilities of a Docker runtime environment (client).

Shorthand Name	Container.Docker
Type Qualified Name	tosca:Container.Docker
Type URI	tosca.capabilities.Container.Docker

3679 **9.2.1.1 Properties**

Name	Required	Type	Constraints	Description
version	no	version[]	None	The Docker version capability (i.e., the versions supported by the capability).
publish_all	no	boolean	default: false	Indicates that all ports (ranges) listed in the <i>dockerfile</i> using the EXPOSE keyword be published.
publish_ports	no	list of PortSpec	None	List of ports mappings from source (Docker container) to target (host) ports to publish.
expose_ports	no	list of PortSpec	None	List of ports mappings from source (Docker container) to expose to other Docker containers (not accessible outside host).
volumes	no	list of string	None	The <i>dockerfile</i> VOLUME command which is used to enable access from the Docker container to a directory on the host machine.
host_id	no	string	None	The optional identifier of an existing host resource that should be used to run this container on.
volume_id	no	string	None	The optional identifier of an existing storage volume (resource) that should be used to create the container's mount point(s) on.

3680 **9.2.1.2 Definition**

```
tosca.capabilities.Container.Docker:
  derived_from: tosca.capabilities.Container
  properties:
    version:
      type: list
      required: false
      entry_schema: version
    publish_all:
      type: boolean
      default: false
      required: false
    publish_ports:
      type: list
      entry_schema: PortSpec
      required: false
    expose_ports:
      type: list
      entry_schema: PortSpec
      required: false
    volumes:
      type: list
      entry_schema: string
      required: false
```

3681 **9.2.1.3 Notes**

- 3682 • When the **expose_ports** property is used, only the **source** and **source_range** properties of
3683 **PortSpec** would be valid for supplying port numbers or ranges, the **target** and **target_range**
3684 properties would be ignored.

3685 **9.3 Node Types**

3686 This section contains non-normative node types referenced in use cases and examples. All additional
3687 Attributes, Properties, Requirements and Capabilities shown in their definitions (and are not inherited
3688 from ancestor normative types) are also considered to be non-normative.

3689 **9.3.1 tosca.nodes.Database.MySQL**

3690 **9.3.1.1 Properties**

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

3691 **9.3.1.2 Definition**

```
tosca.nodes.Database.MySQL:  
  derived_from: tosca.nodes.Database  
  requirements:  
    - host:  
      node: tosca.nodes.DBMS.MySQL
```

3692 **9.3.2 tosca.nodes.DBMS.MySQL**

3693 **9.3.2.1 Properties**

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

3694 **9.3.2.2 Definition**

```
tosca.nodes.DBMS.MySQL:  
  derived_from: tosca.nodes.DBMS  
  properties:  
    port:  
      description: reflect the default MySQL server port  
      default: 3306  
    root_password:  
      # MySQL requires a root_password for configuration  
      # Override parent DBMS definition to make this property required  
      required: true  
  capabilities:  
    # Further constrain the 'host' capability to only allow MySQL databases  
    host:  
      valid_source_types: [ tosca.nodes.Database.MySQL ]
```

3695 9.3.3 `tosca.nodes.WebServer.Apache`

3696 9.3.3.1 Properties

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

3697 9.3.3.2 Definition

```
tosca.nodes.WebServer.Apache:  
  derived_from: tosca.nodes.WebServer
```

3698 9.3.4 `tosca.nodes.WebApplication.WordPress`

3699 This section defines a non-normative Node type for the WordPress [[WordPress](#)] application.

3700 9.3.4.1 Properties

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

3701 9.3.4.2 Definition

```
tosca.nodes.WebApplication.WordPress:  
  derived_from: tosca.nodes.WebApplication  
  properties:  
    admin_user:  
      type: string  
    admin_password:  
      type: string  
    db_host:  
      type: string  
  requirements:  
    - database_endpoint:  
      capability: tosca.capabilities.Endpoint.Database  
      node: tosca.nodes.Database  
      relationship: tosca.relationships.ConnectsTo
```

3702 9.3.5 `tosca.nodes.WebServer.Nodejs`

3703 This non-normative node type represents a Node.js [[NodeJS](#)] web application server.

3704 9.3.5.1 Properties

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

3705 9.3.5.2 Definition

```
tosca.nodes.WebServer.Nodejs:  
  derived_from: tosca.nodes.WebServer  
  properties:  
    # Property to supply the desired implementation in the Github repository
```

```

github_url:
  required: no
  type: string
  description: location of the application on the github.
  default: https://github.com/mmm/testnode.git
interfaces:
  Standard:
    inputs:
      github_url:
        type: string

```

3706 9.3.6 **tosca.nodes.Container.Application.Docker**

3707 9.3.6.1 **Properties**

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

3708 9.3.6.2 **Definition**

```

tosca.nodes.Container.Application.Docker:
  derived_from:
    tosca.nodes.ContainerApplication
  requirements:
    - host:
        capability: tosca.capabilities.Container.Docker

```

3709

10 Component Modeling Use Cases

3710 This section is **non-normative** and includes use cases that explore how to model components and their
3711 relationships using TOSCA Simple Profile in YAML.

3712 10.1.1 Use Case: Exploring the HostedOn relationship using 3713 WebApplication and WebServer

3714 This use case examines the ways TOSCA YAML can be used to express a simple hosting relationship
3715 (i.e., **HostedOn**) using the normative TOSCA **WebServer** and **WebApplication** node types defined in this
3716 specification.

3717 10.1.1.1 WebServer declares its “host” capability

3718 For convenience, relevant parts of the normative TOSCA Node Type for **WebServer** are shown below:

```
tosca.nodes.WebServer
  derived_from: SoftwareComponent
  capabilities:
    ...
  host:
    type: toska.capabilities.Container
    valid_source_types: [ toska.nodes.WebApplication ]
```

3719 As can be seen, the **WebServer** Node Type declares its capability to “contain” (i.e., host) other nodes
3720 using the symbolic name “**host**” and providing the Capability Type **toska.capabilities.Container**. It
3721 should be noted that the symbolic name of “**host**” is not a reserved word, but one assigned by the type
3722 designer that implies at or betokens the associated capability. The **Container** capability definition also
3723 includes a required list of valid Node Types that can be contained by this, the **WebServer**, Node Type.
3724 This list is declared using the keyname of **valid_source_types** and in this case it includes only allowed
3725 type **WebApplication**.

3726 10.1.1.2 WebApplication declares its “host” requirement

3727 The **WebApplication** node type needs to be able to describe the type of capability a target node would
3728 have to provide in order to “host” it. The normative TOSCA capability type **toska.capabilities.Container** is
3729 used to describe all normative TOSCA hosting (i.e., container-containee pattern) relationships. As can be
3730 seen below, the **WebApplication** accomplishes this by declaring a requirement with the symbolic name
3731 “**host**” with the **capability** keyname set to **toska.capabilities.Container**.

3732 Again, for convenience, the relevant parts of the normative **WebApplication** Node Type are shown below:

```
tosca.nodes.WebApplication:
  derived_from: toska.nodes.Root
  requirements:
    - host:
      capability: toska.capabilities.Container
      node: toska.nodes.WebServer
      relationship: toska.relationships.HostedOn
```


3733 **10.1.1.2.1 Notes**

- 3734
- The symbolic name “host” is not a keyword and was selected for consistent use in TOSCA normative node types to give the reader an indication of the type of requirement being referenced. A valid HostedOn relationship could still be established between WebApplicaton and WebServer in a TOSCA Service Template regardless of the symbolic name assigned to either the requirement or capability declaration.
- 3735
- 3736
- 3737
- 3738

3739 **10.1.2 Use Case: Establishing a ConnectsTo relationship to WebServer**

3740 This use case examines the ways TOSCA YAML can be used to express a simple connection relationship (i.e., [ConnectsTo](#)) between some service derived from the [SoftwareComponent](#) Node Type, to the normative [WebServer](#) node type defined in this specification.

3743 The service template that would establish a [ConnectsTo](#) relationship as follows:

```
node_types:
  MyServiceType:
    derived_from: SoftwareComponent
    requirements:
      # This type of service requires a connection to a WebServer's data_endpoint
      - connection1:
          node: WebServer
          relationship: ConnectsTo
          capability: Endpoint

topology_template:
  node_templates:
    my_web_service:
      type: MyServiceType
      ...
      requirements:
        - connection1:
            node: my_web_server

    my_web_server:
      # Note, the normative WebServer node type declares the "data_endpoint"
      # capability of type tosca.capabilities.Endpoint.
      type: WebServer
```

3744 Since the normative **WebServer** Node Type only declares one capability of type
3745 **tosca.capabilities.Endpoint** (or **Endpoint**, its shortname alias in TOSCA) using the symbolic name
3746 **data_endpoint**, the **my_web_service** node template does not need to declare that symbolic name on its
3747 requirement declaration. If however, the **my_web_server** node was based upon some other node type
3748 that declared more than one capability of type **Endpoint**, then the **capability** keyname could be used
3749 to supply the desired symbolic name if necessary.

3750 **10.1.2.1 Best practice**

3751 It should be noted that the best practice for designing Node Types in TOSCA should not export two
3752 capabilities of the same type if they truly offer different functionality (i.e., different capabilities) which
3753 should be distinguished using different Capability Type definitions.

3754 **10.1.3 Use Case: Attaching (local) BlockStorage to a Compute node**

3755 This use case examines the ways TOSCA YAML can be used to express a simple AttachesTo
3756 relationship between a Compute node and a locally attached BlockStorage node.

3757 The service template that would establish an [AttachesTo](#) relationship follows:

```
node_templates:
  my_server:
    type: Compute
    ...
    requirements:
      # contextually this can only be a relationship type
      - local_storage:
          # capability is provided by Compute Node Type
          node: my_block_storage
          relationship:
            type: AttachesTo
            properties:
              location: /path1/path2
          # This maps the local requirement name 'local_storage' to the
          # target node's capability name 'attachment'

  my_block_storage:
    type: BlockStorage
    properties:
      size: 10 GB
```

3758 **10.1.4 Use Case: Reusing a BlockStorage Relationship using Relationship**
3759 **Type or Relationship Template**

3760 This builds upon the previous use case (10.1.3) to examine how a template author could attach multiple
3761 Compute nodes (templates) to the same BlockStorage node (template), but with slightly different property
3762 values for the AttachesTo relationship.

3763

3764 Specifically, several notation options are shown (in this use case) that achieve the same desired result.

3765 **10.1.4.1 Simple Profile Rationale**

3766 Referencing an explicitly declared Relationship Template is a convenience of the Simple Profile that
3767 allows template authors an entity to set, constrain or override the properties and operations as defined in
3768 its declared (Relationship) Type much as allowed now for Node Templates. It is especially useful when a
3769 complex Relationship Type (with many configurable properties or operations) has several logical
3770 occurrences in the same Service (Topology) Template; allowing the author to avoid configuring these
3771 same properties and operations in multiple Node Templates.

3772 **10.1.4.2 Notation Style #1: Augment AttachesTo Relationship Type directly in**
3773 **each Node Template**

3774 This notation extends the methodology used for establishing a HostedOn relationship, but allowing
3775 template author to supply (dynamic) configuration and/or override of properties and operations.

3776
3777 **Note:** This option will remain valid for Simple Profile regardless of other notation (copy or aliasing) options
3778 being discussed or adopted for future versions.

3779

```
node_templates:

  my_block_storage:
    type: BlockStorage
    properties:
      size: 10

  my_web_app_tier_1:
    type: Compute
    requirements:
      - local_storage:
          node: my_block_storage
          relationship: MyAttachesTo
          # use default property settings in the Relationship Type definition

  my_web_app_tier_2:
    type: Compute
    requirements:
      - local_storage:
          node: my_block_storage
          relationship:
            type: MyAttachesTo
            # Override default property setting for just the 'location' property
            properties:
              location: /some_other_data_location

relationship_types:

  MyAttachesTo:
    derived_from: AttachesTo
    properties:
      location: /default_location
    interfaces:
      Configure:
        post_configure_target:
          implementation: default_script.sh
```

3780

3781 **10.1.4.3 Notation Style #2: Use the ‘template’ keyword on the Node Templates to**
3782 **specify which named Relationship Template to use**

3783 This option shows how to explicitly declare different named Relationship Templates within the Service
3784 Template as part of a **relationship_templates** section (which have different property values) and can
3785 be referenced by different Compute typed Node Templates.

3786

```
node_templates:

  my_block_storage:
    type: BlockStorage
    properties:
      size: 10

  my_web_app_tier_1:
    derived_from: Compute
    requirements:
      - local_storage:
          node: my_block_storage
          relationship: storage_attachesto_1

  my_web_app_tier_2:
    derived_from: Compute
    requirements:
      - local_storage:
          node: my_block_storage
          relationship: storage_attachesto_2

relationship_templates:

  storage_attachesto_1:
    type: MyAttachesTo
    properties:
      location: /my_data_location

  storage_attachesto_2:
    type: MyAttachesTo
    properties:
      location: /some_other_data_location

relationship_types:

  MyAttachesTo:
    derived_from: AttachesTo
    interfaces:
      some_interface_name:
        some_operation:
          implementation: default_script.sh
```

3787

3788 **10.1.4.4 Notation Style #3: Using the “copy” keyname to define a similar**
3789 **Relationship Template**

3790 How does TOSCA make it easier to create a new relationship template that is mostly the same as one
3791 that exists without manually copying all the same information? TOSCA provides the **copy** keyname as a
3792 convenient way to copy an existing template definition into a new template definition as a starting point or
3793 basis for describing a new definition and avoid manual copy. The end results are cleaner TOSCA Service
3794 Templates that allows the description of only the changes (or deltas) between similar templates.

3795 The example below shows that the Relationship Template named **storage_attachesto_1** provides
3796 some overrides (conceptually a large set of overrides) on its Type which the Relationship Template
3797 named **storage_attachesto_2** wants to “**copy**” before perhaps providing a smaller number of overrides.

```
node_templates:

  my_block_storage:
    type: BlockStorage
    properties:
      size: 10

  my_web_app_tier_1:
    derived_from: Compute
    requirements:
      - attachment:
          node: my_block_storage
          relationship: storage_attachesto_1

  my_web_app_tier_2:
    derived_from: Compute
    requirements:
      - attachment:
          node: my_block_storage
          relationship: storage_attachesto_2

relationship_templates:

  storage_attachesto_1:
    type: MyAttachesTo
    properties:
      location: /my_data_location
    interfaces:
      some_interface_name:
        some_operation_name_1: my_script_1.sh
        some_operation_name_2: my_script_2.sh
        some_operation_name_3: my_script_3.sh

  storage_attachesto_2:
    # Copy the contents of the “storage_attachesto_1” template into this new one
    copy: storage_attachesto_1
```

```
# Then change just the value of the location property
properties:
  location: /some_other_data_location

relationship_types:

MyAttachesTo:
  derived_from: AttachesTo
  interfaces:
    some_interface_name:
      some_operation:
        implementation: default_script.sh
```

3798

11 Application Modeling Use Cases

3799 This section is **non-normative** and includes use cases that show how to model Infrastructure-as-a-
 3800 Service (IaaS), Platform-as-a-Service (PaaS) and complete application uses cases using TOSCA Simple
 3801 Profile in YAML.

11.1 Use cases

3803 Many of the use cases listed below can be found under the following link:

3804 <https://github.com/openstack/heat-translator/tree/master/translator/tests/data>

11.1.1 Overview

3805

Name	Description
Compute: Create a single Compute instance with a host Operating System	Introduces a TOSCA Compute node type which is used to stand up a single compute instance with a host Operating System Virtual Machine (VM) image selected by the platform provider using the Compute node's properties.
Software Component 1: Automatic deployment of a Virtual Machine (VM) image artifact	Introduces the SoftwareComponent node type which declares software that is hosted on a Compute instance. In this case, the SoftwareComponent declares a VM image as a deployment artifact which includes its own pre-packaged operating system and software. The TOSCA Orchestrator detects this known deployment artifact type on the SoftwareComponent node template and automatically deploys it to the Compute node.
BlockStorage-1: Attaching Block Storage to a single Compute instance	Demonstrates how to attach a TOSCA BlockStorage node to a Compute node using the normative AttachesTo relationship.
BlockStorage-2: Attaching Block Storage using a custom Relationship Type	Demonstrates how to attach a TOSCA BlockStorage node to a Compute node using a custom RelationshipType that derives from the normative AttachesTo relationship.
BlockStorage-3: Using a Relationship Template of type AttachesTo	Demonstrates how to attach a TOSCA BlockStorage node to a Compute node using a TOSCA Relationship Template that is based upon the normative AttachesTo Relationship Type.
BlockStorage-4: Single Block Storage shared by 2-Tier Application with custom AttachesTo Type and implied relationships	This use case shows 2 Compute instances (2 tiers) with one BlockStorage node, and also uses a custom AttachesTo Relationship that provides a default mount point (i.e., location) which the 1 st tier uses, but the 2 nd tier provides a different mount point.
BlockStorage-5: Single Block Storage shared by 2-Tier Application with custom AttachesTo Type and explicit Relationship Templates	This use case is like the previous BlockStorage-4 use case, but also creates two relationship templates (one for each tier) each of which provide a different mount point (i.e., location) which overrides the default location defined in the custom Relationship Type.
BlockStorage-6: Multiple Block Storage attached to different Servers	This use case demonstrates how two different TOSCA BlockStorage nodes can be attached to two different Compute nodes (i.e., servers) each using the normative AttachesTo relationship.
Object Storage 1: Creating an Object Storage service	Introduces the TOSCA ObjectStorage node type and shows how it can be instantiated.
Network-1: Server bound to a new network	Introduces the TOSCA Network and Port nodes used for modeling logical networks using the LinksTo and BindsTo Relationship Types. In this use case, the template is invoked without an existing network_name as an input property so a new network is created using the properties declared in the Network node.
Network-2: Server	Shows how to use a network_name as an input parameter to the template to allow a server to be

bound to an existing network	associated with (i.e. bound to) an existing Network .
Network-3: Two servers bound to a single network	This use case shows how two servers (Compute nodes) can be associated with the same Network node using two logical network Ports .
Network-4: Server bound to three networks	This use case shows how three logical networks (Network nodes), each with its own IP address range, can be associated with the same server (Compute node).
WebServer-DBMS-1: WordPress [WordPress] + MySQL, single instance	Shows how to host a TOSCA WebServer with a TOSCA WebApplication , DBMS and Database Node Types along with their dependent HostedOn and ConnectsTo relationships.
WebServer-DBMS-2: Nodejs with PayPal Sample App and MongoDB on separate instances	Instantiates a 2-tier application with Nodejs and its (PayPal sample) WebApplication on one tier which connects a MongoDB database (which stores its application data) using a ConnectsTo relationship.
Multi-Tier-1: Elasticsearch, Logstash, Kibana (ELK)	Shows Elasticsearch , Logstash and Kibana (ELK) being used in a typical manner to collect, search and monitor/visualize data from a running application. This use case builds upon the previous Nodejs/MongoDB 2-tier application as the one being monitored. The collectd and rsyslog components are added to both the WebServer and Database tiers which work to collect data for Logstash. In addition to the application tiers, a 3 rd tier is introduced with Logstash to collect data from the application tiers. Finally a 4 th tier is added to search the Logstash data with Elasticsearch and visualize it using Kibana . Note: This use case also shows the convenience of using a single YAML macro (declared in the dsl_definitions section of the TOSCA Service Template) on multiple Compute nodes.
Container-1: Containers using Docker single Compute instance (Containers only)	Minimalist TOSCA Service Template description of 2 Docker containers linked to each other. Specifically, one container runs wordpress and connects to second mysql database container both on a single server (i.e., Compute instance). The use case also demonstrates how TOSCA declares and references Docker images from the Docker Hub repository. Variation 1: Docker Container nodes (only) providing their Docker Requirements allowing platform (orchestrator) to select/provide the underlying Docker implementation (Capability).

3806 11.1.2 Compute: Create a single Compute instance with a host Operating System

3808 11.1.2.1 Description

3809 This use case demonstrates how the TOSCA Simple Profile specification can be used to stand up a
3810 single Compute instance with a guest Operating System using a normative TOSCA **Compute** node. The
3811 TOSCA Compute node is declarative in that the service template describes both the processor and host
3812 operating system platform characteristics (i.e., properties declared on the capability named “**os**”
3813 sometimes called a “flavor”) that are desired by the template author. The cloud provider would attempt to
3814 fulfill these properties (to the best of its abilities) during orchestration.

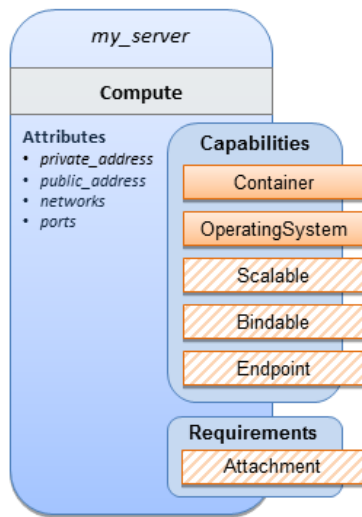
3815 11.1.2.2 Features

3816 This use case introduces the following TOSCA Simple Profile features:

- 3817 • A node template that uses the normative TOSCA **Compute** Node Type along with showing an
3818 exemplary set of its properties being configured.

- 3819 • Use of the TOSCA Service Template **inputs** section to declare a configurable value the template
- 3820 user may supply at runtime. In this case, the “host” property named “num_cpus” (of type integer)
- 3821 is declared.
 - 3822 ○ Use of a property constraint to limit the allowed integer values for the “num_cpus”
 - 3823 property to a specific list supplied in the property declaration.
- 3824 • Use of the TOSCA Service Template **outputs** section to declare a value the template user may
- 3825 request at runtime. In this case, the property named “instance_ip” is declared
 - 3826 ○ The “instance_ip” output property is programmatically retrieved from the **Compute**
 - 3827 node’s “public_address” attribute using the TOSCA Service Template-level
 - 3828 **get_attribute** function.

3829 11.1.2.3 Logical Diagram



3830

3831 11.1.2.4 Sample YAML

```

tosca_definitions_version: tosca_simple_yaml_1_0

description: >
  TOSCA simple profile that just defines a single compute instance and selects a
  (guest) host Operating System from the Compute node’s properties. Note, this
  example does not include default values on inputs properties.

topology_template:
  inputs:
    cpus:
      type: integer
      description: Number of CPUs for the server.
      constraints:
        - valid_values: [ 1, 2, 4, 8 ]

  node_templates:
    my_server:
      type: Compute
      capabilities:
        host:
      properties:
  
```

```

    disk_size: 10 GB
    num_cpus: { get_input: cpus }
    mem_size: 1 GB
  os:
    properties:
      architecture: x86_64
      type: Linux
      distribution: ubuntu
      version: 12.04
  outputs:
    private_ip:
      description: The private IP address of the deployed server instance.
      value: { get_attribute: [my_server, private_address] }

```

3832 11.1.2.5 Notes

- 3833 • This use case uses a versioned, Linux Ubuntu distribution on the Compute node.

3834 11.1.3 Software Component 1: Automatic deployment of a Virtual Machine 3835 (VM) image artifact

3836 11.1.3.1 Description

3837 This use case demonstrates how the TOSCA SoftwareComponent node type can be used to declare
3838 software that is packaged in a standard Virtual Machine (VM) image file format (i.e., in this case QCOW2)
3839 and is hosted on a TOSCA Compute node (instance). In this variation, the SoftwareComponent declares
3840 a VM image as a deployment artifact that includes its own pre-packaged operating system and software.
3841 The TOSCA Orchestrator detects this known deployment artifact type on the SoftwareComponent node
3842 template and automatically deploys it to the Compute node.

3843 11.1.3.2 Features

3844 This use case introduces the following TOSCA Simple Profile features:

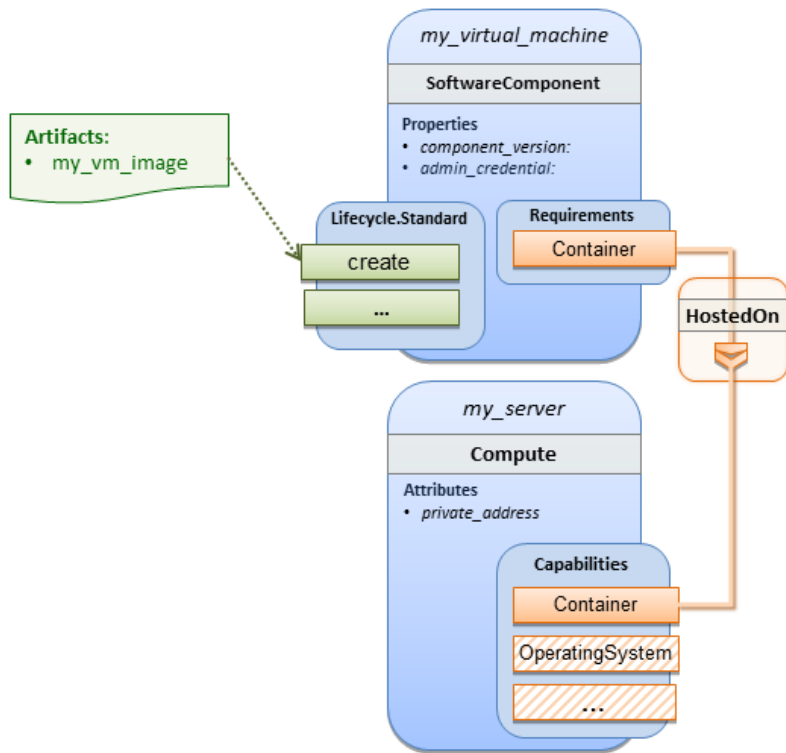
- 3845 • A node template that uses the normative TOSCA **SoftwareComponent** Node Type along with
3846 showing an exemplary set of its properties being configured.
- 3847 • Use of the TOSCA Service Template **artifacts** section to declare a Virtual Machine (VM) image
3848 artifact type which is referenced by the **SoftwareComponent** node template.
- 3849 • The VM file format, in this case QCOW2, includes its own guest Operating System (OS) and
3850 therefore does **not** “require” a TOSCA **OperatingSystem** capability from the TOSCA Compute
3851 node.

3852 11.1.3.3 Assumptions

3853 This use case assumes the following:

- 3854 • That the TOSCA Orchestrator (working with the Cloud provider’s underlying management
3855 services) is able to instantiate a Compute node that has a hypervisor that supports the Virtual
3856 Machine (VM) image format, in this case QCOW2, which should be compatible with many
3857 standard hypervisors such as XEN and KVM.
- 3858 • This is not a “bare metal” use case and assumes the existence of a hypervisor on the machine
3859 that is allocated to “host” the Compute instance supports (e.g. has drivers, etc.) the VM image
3860 format in this example.

3861 **11.1.3.4 Logical Diagram**



3862

3863 **11.1.3.5 Sample YAML**

```

tosca_definitions_version: toska_simple_yaml_1_0

description: >
  TOSCA Simple Profile with a SoftwareComponent node with a declared Virtual
  machine (VM) deployment artifact that automatically deploys to its host Compute
  node.

topology_template:

  node_templates:
    my_virtual_machine:
      type: SoftwareComponent
      artifacts:
        my_vm_image:
          file: images/fedora-18-x86_64.qcow2
          type: toska.artifacts.Deployment.Image.VM.QCOW2
      requirements:
        - host: my_server
        # Automatically deploy the VM image referenced on the create operation
      interfaces:
        Standard:
          create: my_vm_image

        # Compute instance with no Operating System guest host
    my_server:
      type: Compute
  
```

```

capabilities:
  # Note: no guest OperatingSystem requirements as these are in the image.
  host:
    properties:
      disk_size: 10 GB
      num_cpus: { get_input: cpus }
      mem_size: 4 GB

outputs:
  private_ip:
    description: The private IP address of the deployed server instance.
    value: { get_attribute: [my_server, private_address] }

```

3864 **11.1.3.6 Notes**

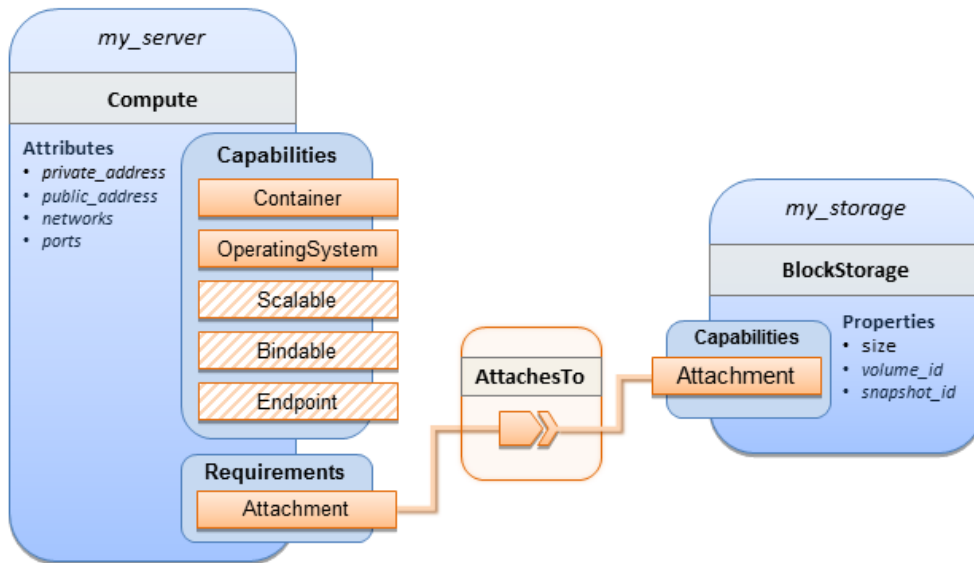
- 3865 • The use of the **type** keyname on the **artifact** definition (within the **my_virtual_machine** node
3866 template) to declare the ISO image deployment artifact type (i.e.,
3867 **tosca.artifacts.Deployment.Image.VM.ISO**) is redundant since the file extension is “.iso”
3868 which associated with this known, declared artifact type.
- 3869 • This use case references a filename on the **my_vm_image** artifact, which indicates a Linux,
3870 Fedora 18, x86 VM image, only as one possible example.

3871 **11.1.4 Block Storage 1: Using the normative AttachesTo Relationship Type**

3872 **11.1.4.1 Description**

3873 This use case demonstrates how to attach a TOSCA **BlockStorage** node to a **Compute** node using the
3874 normative **AttachesTo** relationship.

3875 **11.1.4.2 Logical Diagram**



3876

3877 **11.1.4.3 Sample YAML**

```
tosca_definitions_version: tosca_simple_yaml_1_0
```

```

description: >
  TOSCA simple profile with server and attached block storage using the normative
  AttachesTo Relationship Type.

topology_template:

  inputs:
    cpus:
      type: integer
      description: Number of CPUs for the server.
      constraints:
        - valid_values: [ 1, 2, 4, 8 ]
    storage_size:
      type: scalar-unit.size
      description: Size of the storage to be created.
      default: 1 GB
    storage_snapshot_id:
      type: string
      description: >
        Optional identifier for an existing snapshot to use when creating storage.
    storage_location:
      type: string
      description: Block storage mount point (filesystem path).

  node_templates:
    my_server:
      type: Compute
      capabilities:
        host:
          properties:
            disk_size: 10 GB
            num_cpus: { get_input: cpus }
            mem_size: 1 GB
        os:
          properties:
            architecture: x86_64
            type: linux
            distribution: fedora
            version: 18.0
      requirements:
        - local_storage:
            node: my_storage
            relationship:
              type: AttachesTo
              properties:
                location: { get_input: storage_location }

    my_storage:
      type: BlockStorage
      properties:
        size: { get_input: storage_size }
        snapshot_id: { get_input: storage_snapshot_id }

  outputs:

```

```

private_ip:
  description: The private IP address of the newly created compute instance.
  value: { get_attribute: [my_server, private_address] }
volume_id:
  description: The volume id of the block storage instance.
  value: { get_attribute: [my_storage, volume_id] }

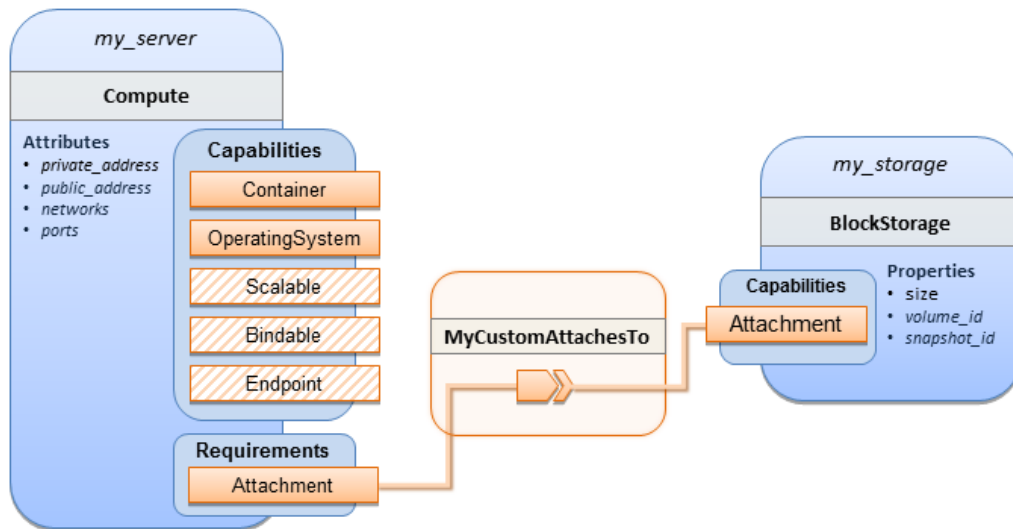
```

3878 11.1.5 Block Storage 2: Using a custom AttachesTo Relationship Type

3879 11.1.5.1 Description

3880 This use case demonstrates how to attach a TOSCA **BlockStorage** node to a **Compute** node using a
 3881 custom RelationshipType that derives from the normative **AttachesTo** relationship.

3882 11.1.5.2 Logical Diagram



3883

3884 11.1.5.3 Sample YAML

3885

```

tosca_definitions_version: tosca_simple_yaml_1_0

description: >
  TOSCA simple profile with server and attached block storage using a custom
  AttachesTo Relationship Type.

relationship_types:
  MyCustomAttachesTo:
    derived_from: AttachesTo

topology_template:
  inputs:
    cpus:
      type: integer
      description: Number of CPUs for the server.
      constraints:
        - valid_values: [ 1, 2, 4, 8 ]

```

```

storage_size:
  type: scalar-unit.size
  description: Size of the storage to be created.
  default: 1 GB
storage_snapshot_id:
  type: string
  description: >
    Optional identifier for an existing snapshot to use when creating storage.
storage_location:
  type: string
  description: Block storage mount point (filesystem path).

node_templates:
  my_server:
    type: Compute
    capabilities:
      host:
        properties:
          disk_size: 10 GB
          num_cpus: { get_input: cpus }
          mem_size: 4 GB
        os:
          properties:
            architecture: x86_64
            type: Linux
            distribution: Fedora
            version: 18.0
      requirements:
        - local_storage:
            node: my_storage
            # Declare custom AttachesTo type using the 'relationship' keyword
            relationship:
              type: MyCustomAttachesTo
              properties:
                location: { get_input: storage_location }
    my_storage:
      type: BlockStorage
      properties:
        size: { get_input: storage_size }
        snapshot_id: { get_input: storage_snapshot_id }

outputs:
  private_ip:
    description: The private IP address of the newly created compute instance.
    value: { get_attribute: [my_server, private_address] }
  volume_id:
    description: The volume id of the block storage instance.
    value: { get_attribute: [my_storage, volume_id] }

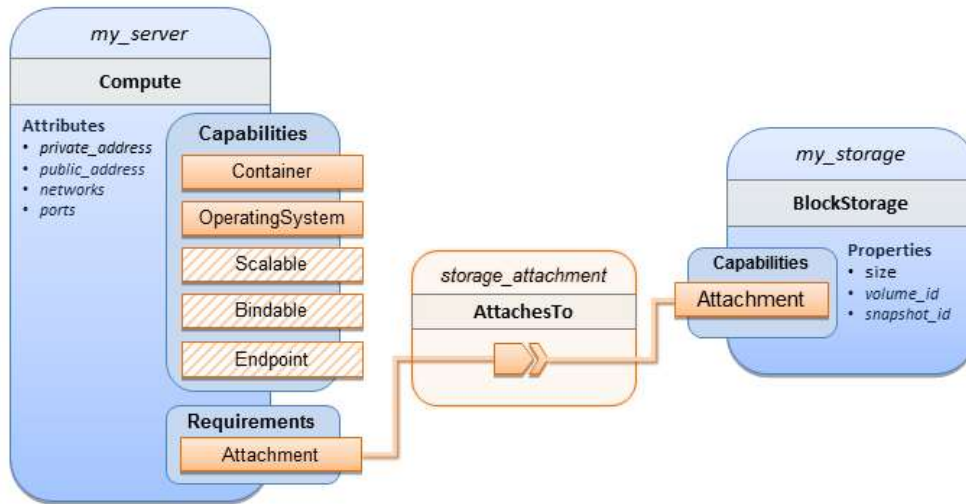
```

3886 11.1.6 Block Storage 3: Using a Relationship Template of type AttachesTo

3887 11.1.6.1 Description

3888 This use case demonstrates how to attach a TOSCA **BlockStorage** node to a **Compute** node using a
 3889 TOSCA Relationship Template that is based upon the normative **AttachesTo** Relationship Type.

3890 **11.1.6.2 Logical Diagram**



3891

3892 **11.1.6.3 Sample YAML**

3893

```

tosca_definitions_version: tosca_simple_yaml_1_0

description: >
  TOSCA simple profile with server and attached block storage using a named
  Relationship Template for the storage attachment.

topology_template:
  inputs:
    cpus:
      type: integer
      description: Number of CPUs for the server.
      constraints:
        - valid_values: [ 1, 2, 4, 8 ]
    storage_size:
      type: scalar-unit.size
      description: Size of the storage to be created.
      default: 1 GB
    storage_location:
      type: string
      description: Block storage mount point (filesystem path).

  node_templates:
    my_server:
      type: Compute
      capabilities:
        host:
          properties:
            disk_size: 10 GB
            num_cpus: { get_input: cpus }
            mem_size: 4 GB
      os:
        properties:
          architecture: x86_64
  
```



```

    type: Linux
    distribution: Fedora
    version: 18.0
  requirements:
    - local_storage:
        node: my_storage
        # Declare template to use with 'relationship' keyword
        relationship: storage_attachment

  my_storage:
    type: BlockStorage
    properties:
      size: { get_input: storage_size }

  relationship_templates:
    storage_attachment:
      type: AttachesTo
      properties:
        location: { get_input: storage_location }

  outputs:
    private_ip:
      description: The private IP address of the newly created compute instance.
      value: { get_attribute: [my_server, private_address] }
    volume_id:
      description: The volume id of the block storage instance.
      value: { get_attribute: [my_storage, volume_id] }

```

3894 **11.1.7 Block Storage 4: Single Block Storage shared by 2-Tier Application**
 3895 **with custom AttachesTo Type and implied relationships**

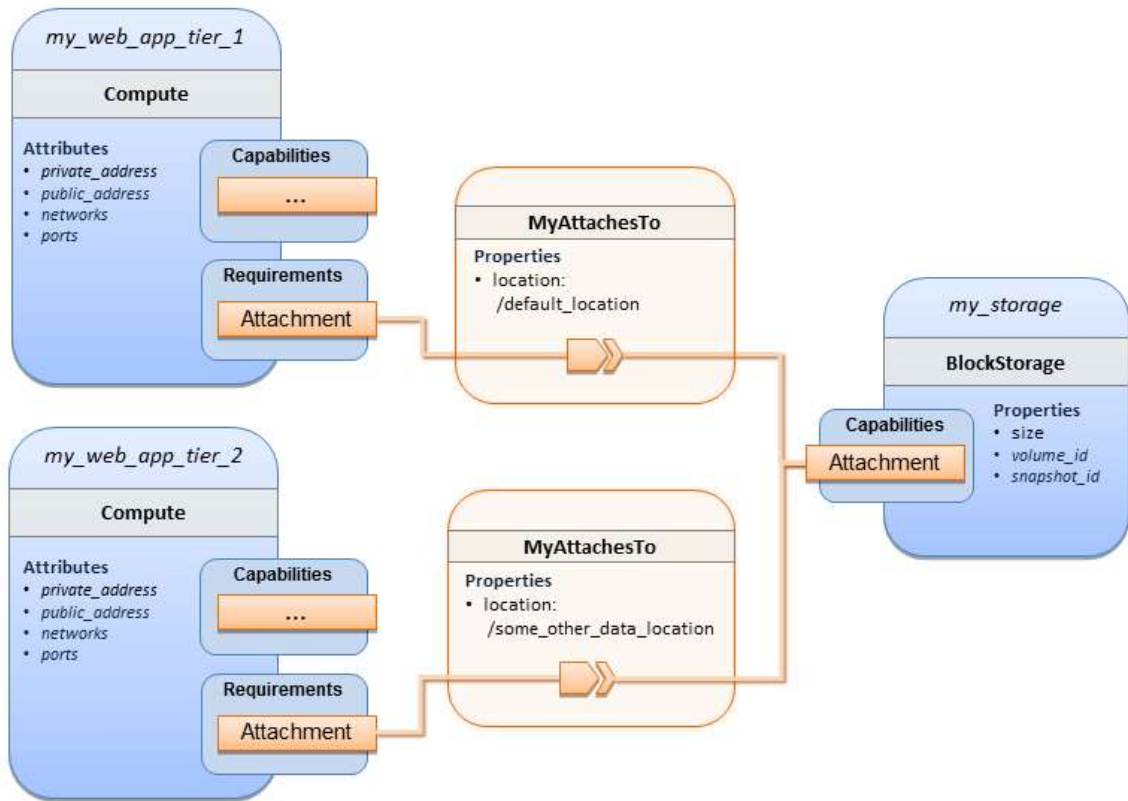
3896 **11.1.7.1 Description**

3897 This use case shows 2 compute instances (2 tiers) with one BlockStorage node, and also uses a custom
 3898 **AttachesTo** Relationship that provides a default mount point (i.e., **location**) which the 1st tier uses,
 3899 but the 2nd tier provides a different mount point.

3900

3901 Please note that this use case assumes both Compute nodes are accessing different directories within
 3902 the shared, block storage node to avoid collisions.

3903 **11.1.7.2 Logical Diagram**



3904

3905 **11.1.7.3 Sample YAML**

```

tosca_definitions_version: tosca_simple_yaml_1_0

description: >
  TOSCA simple profile with a Single Block Storage node shared by 2-Tier Application with
  custom AttachesTo Type and implied relationships.

relationship_types:
  MyAttachesTo:
    derived_from: tosca.relationships.AttachesTo
    properties:
      location:
        type: string
        default: /default_location

topology_template:
  inputs:
    cpus:
      type: integer
      description: Number of CPUs for the server.
      constraints:
        - valid_values: [ 1, 2, 4, 8 ]
    storage_size:
      type: scalar-unit.size
  
```

```

    default: 1 GB
    description: Size of the storage to be created.
storage_snapshot_id:
  type: string
  description: >
    Optional identifier for an existing snapshot to use when creating storage.

node_templates:
  my_web_app_tier_1:
    type: tosca.nodes.Compute
    capabilities:
      host:
        properties:
          disk_size: 10 GB
          num_cpus: { get_input: cpus }
          mem_size: 4096 MB
        os:
          properties:
            architecture: x86_64
            type: Linux
            distribution: Fedora
            version: 18.0
      requirements:
        - local_storage:
            node: my_storage
            relationship: MyAttachesTo

  my_web_app_tier_2:
    type: tosca.nodes.Compute
    capabilities:
      host:
        properties:
          disk_size: 10 GB
          num_cpus: { get_input: cpus }
          mem_size: 4096 MB
        os:
          properties:
            architecture: x86_64
            type: Linux
            distribution: Fedora
            version: 18.0
      requirements:
        - local_storage:
            node: my_storage
            relationship:
              type: MyAttachesTo
              properties:
                location: /some_other_data_location

  my_storage:
    type: tosca.nodes.BlockStorage
    properties:
      size: { get_input: storage_size }
      snapshot_id: { get_input: storage_snapshot_id }

```

```

outputs:
  private_ip_1:
    description: The private IP address of the application's first tier.
    value: { get_attribute: [my_web_app_tier_1, private_address] }
  private_ip_2:
    description: The private IP address of the application's second tier.
    value: { get_attribute: [my_web_app_tier_2, private_address] }
  volume_id:
    description: The volume id of the block storage instance.
    value: { get_attribute: [my_storage, volume_id] }

```

3906 **11.1.8 Block Storage 5: Single Block Storage shared by 2-Tier Application**
 3907 **with custom AttachesTo Type and explicit Relationship Templates**

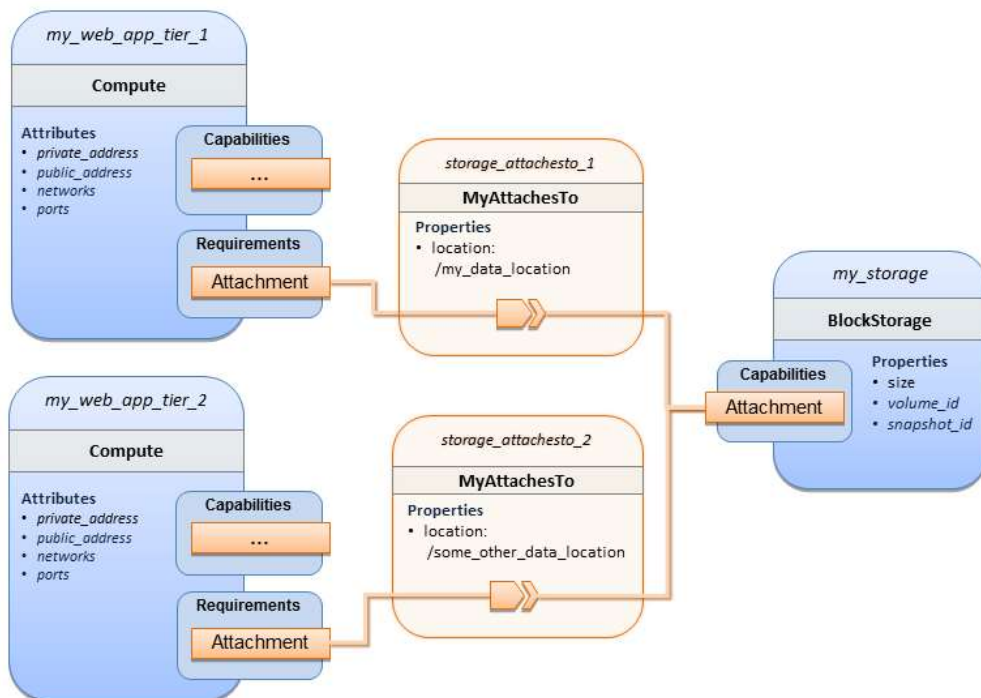
3908 **11.1.8.1 Description**

3909 This use case is like the Notation1 use case, but also creates two relationship templates (one for each
 3910 tier) each of which provide a different mount point (i.e., **location**) which overrides the default location
 3911 defined in the custom Relationship Type.

3912

3913 Please note that this use case assumes both Compute nodes are accessing different directories within
 3914 the shared, block storage node to avoid collisions.

3915 **11.1.8.2 Logical Diagram**



3916

3917 **11.1.8.3 Sample YAML**

```

tosca_definitions_version: tosca_simple_yaml_1_0

description: >

```

TOSCA simple profile with a single Block Storage node shared by 2-Tier Application with custom AttachesTo Type and explicit Relationship Templates.

```
relationship_types:
  MyAttachesTo:
    derived_from: tosca.relationships.AttachesTo
    properties:
      location:
        type: string
        default: /default_location

topology_template:
  inputs:
    cpus:
      type: integer
      description: Number of CPUs for the server.
      constraints:
        - valid_values: [ 1, 2, 4, 8 ]
    storage_size:
      type: scalar-unit.size
      default: 1 GB
      description: Size of the storage to be created.
    storage_snapshot_id:
      type: string
      description: >
        Optional identifier for an existing snapshot to use when creating storage.
    storage_location:
      type: string
      description: >
        Block storage mount point (filesystem path).

node_templates:

  my_web_app_tier_1:
    type: tosca.nodes.Compute
    capabilities:
      host:
        properties:
          disk_size: 10 GB
          num_cpus: { get_input: cpus }
          mem_size: 4096 MB
      os:
        properties:
          architecture: x86_64
          type: Linux
          distribution: Fedora
          version: 18.0
    requirements:
      - local_storage:
          node: my_storage
          relationship: storage_attachesto_1

  my_web_app_tier_2:
    type: tosca.nodes.Compute
    capabilities:
```

```

host:
  properties:
    disk_size: 10 GB
    num_cpus: { get_input: cpus }
    mem_size: 4096 MB
  os:
    properties:
      architecture: x86_64
      type: Linux
      distribution: Fedora
      version: 18.0
  requirements:
    - local_storage:
      node: my_storage
      relationship: storage_attachesto_2

my_storage:
  type: tosca.nodes.BlockStorage
  properties:
    size: { get_input: storage_size }
    snapshot_id: { get_input: storage_snapshot_id }

relationship_templates:
  storage_attachesto_1:
    type: MyAttachesTo
    properties:
      location: /my_data_location

  storage_attachesto_2:
    type: MyAttachesTo
    properties:
      location: /some_other_data_location

outputs:
  private_ip_1:
    description: The private IP address of the application's first tier.
    value: { get_attribute: [my_web_app_tier_1, private_address] }
  private_ip_2:
    description: The private IP address of the application's second tier.
    value: { get_attribute: [my_web_app_tier_2, private_address] }
  volume_id:
    description: The volume id of the block storage instance.
    value: { get_attribute: [my_storage, volume_id] }

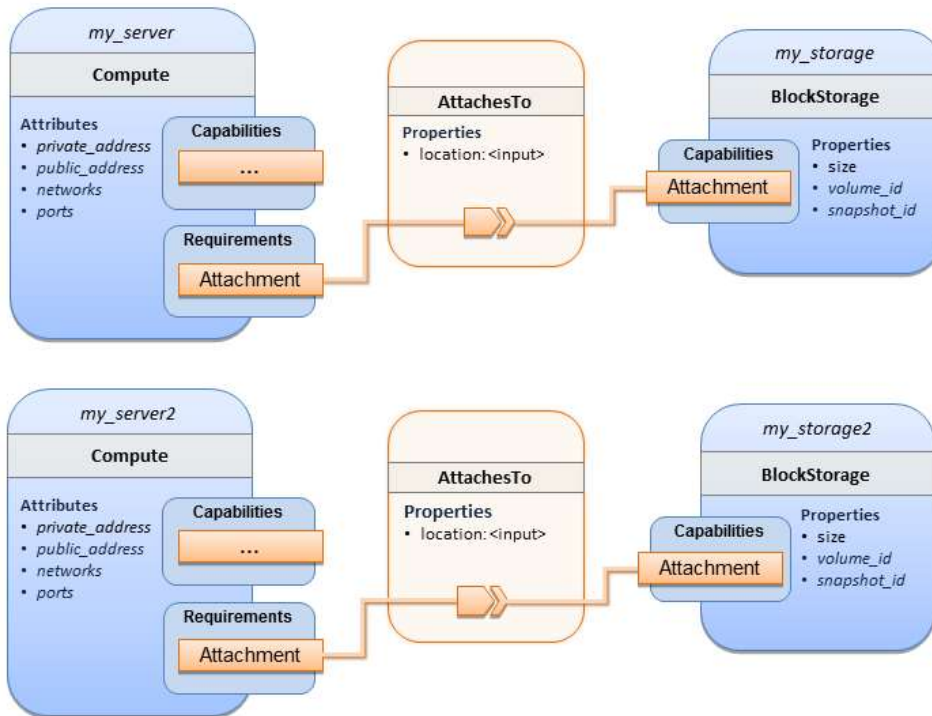
```

3918 11.1.9 Block Storage 6: Multiple Block Storage attached to different Servers

3919 11.1.9.1 Description

3920 This use case demonstrates how two different TOSCA **BlockStorage** nodes can be attached to two
 3921 different **Compute** nodes (i.e., servers) each using the normative **AttachesTo** relationship.

3922 **11.1.9.2 Logical Diagram**



3923

3924 **11.1.9.3 Sample YAML**

```

tosca_definitions_version: tosca_simple_yaml_1_0

description: >
  TOSCA simple profile with 2 servers each with different attached block storage.

topology_template:
  inputs:
    cpus:
      type: integer
      description: Number of CPUs for the server.
      constraints:
        - valid_values: [ 1, 2, 4, 8 ]
    storage_size:
      type: scalar-unit.size
      default: 1 GB
      description: Size of the storage to be created.
    storage_snapshot_id:
      type: string
      description: >
        Optional identifier for an existing snapshot to use when creating storage.
    storage_location:
      type: string
      description: >
        Block storage mount point (filesystem path).

  node_templates:
    my_server:

```

```

type: tosca.nodes.Compute
capabilities:
  host:
    properties:
      disk_size: 10 GB
      num_cpus: { get_input: cpus }
      mem_size: 4096 MB
    os:
      properties:
        architecture: x86_64
        type: Linux
        distribution: Fedora
        version: 18.0
  requirements:
    - local_storage:
        node: my_storage
        relationship:
          type: AttachesTo
          properties:
            location: { get_input: storage_location }
my_storage:
  type: tosca.nodes.BlockStorage
  properties:
    size: { get_input: storage_size }
    snapshot_id: { get_input: storage_snapshot_id }

my_server2:
  type: tosca.nodes.Compute
  capabilities:
    host:
      properties:
        disk_size: 10 GB
        num_cpus: { get_input: cpus }
        mem_size: 4096 MB
      os:
        properties:
          architecture: x86_64
          type: Linux
          distribution: Fedora
          version: 18.0
    requirements:
      - local_storage:
          node: my_storage2
          relationship:
            type: AttachesTo
            properties:
              location: { get_input: storage_location }
my_storage2:
  type: tosca.nodes.BlockStorage
  properties:
    size: { get_input: storage_size }
    snapshot_id: { get_input: storage_snapshot_id }

outputs:
  server_ip_1:

```



```

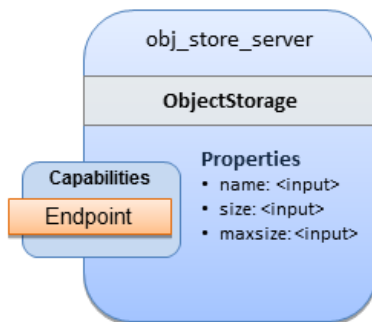
description: The private IP address of the application's first server.
value: { get_attribute: [my_server, private_address] }
server_ip_2:
description: The private IP address of the application's second server.
value: { get_attribute: [my_server2, private_address] }
volume_id_1:
description: The volume id of the first block storage instance.
value: { get_attribute: [my_storage, volume_id] }
volume_id_2:
description: The volume id of the second block storage instance.
value: { get_attribute: [my_storage2, volume_id] }

```

3925 11.1.10 Object Storage 1: Creating an Object Storage service

3926 11.1.10.1 Description

3927 11.1.10.2 Logical Diagram



3928

3929 11.1.10.3 Sample YAML

```

tosca_definitions_version: tosca_simple_yaml_1_0

description: >
  Tosca template for creating an object storage service.

topology_template:
  inputs:
    objectstore_name:
      type: string

  node_templates:
    obj_store_server:
      type: tosca.nodes.ObjectStorage
      properties:
        name: { get_input: objectstore_name }
        size: 4096 MB
        maxsize: 20 GB

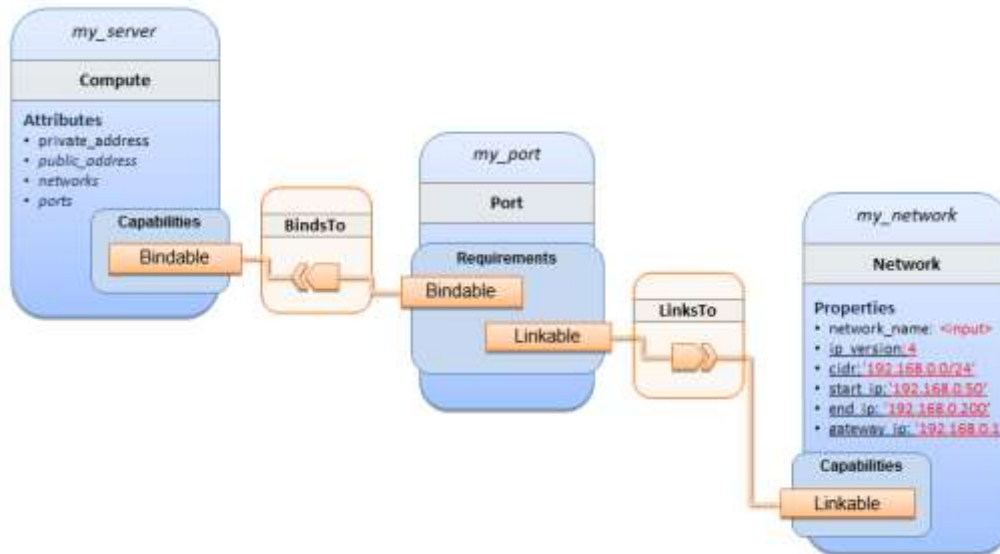
```

3930 **11.1.11 Network 1: Server bound to a new network**

3931 **11.1.11.1 Description**

3932 Introduces the TOSCA **Network** and **Port** nodes used for modeling logical networks using the **LinksTo** and
3933 **BindsTo** Relationship Types. In this use case, the template is invoked without an existing network_name
3934 as an input property so a new network is created using the properties declared in the Network node.

3935 **11.1.11.2 Logical Diagram**



3936

3937 **11.1.11.3 Sample YAML**

```
tosca_definitions_version: tosca_simple_yaml_1_0

description: >
  TOSCA simple profile with 1 server bound to a new network

topology_template:

  inputs:
    network_name:
      type: string
      description: Network name

  node_templates:
    my_server:
      type: tosca.nodes.Compute
      capabilities:
        host:
          properties:
            disk_size: 10 GB
            num_cpus: 1
            mem_size: 4096 MB
      os:
        properties:
          architecture: x86_64
          type: Linux
```

```

distribution: CirrOS
version: 0.3.2

my_network:
  type: toska.nodes.network.Network
  properties:
    network_name: { get_input: network_name }
    ip_version: 4
    cidr: '192.168.0.0/24'
    start_ip: '192.168.0.50'
    end_ip: '192.168.0.200'
    gateway_ip: '192.168.0.1'

my_port:
  type: toska.nodes.network.Port
  requirements:
    - binding: my_server
    - link: my_network

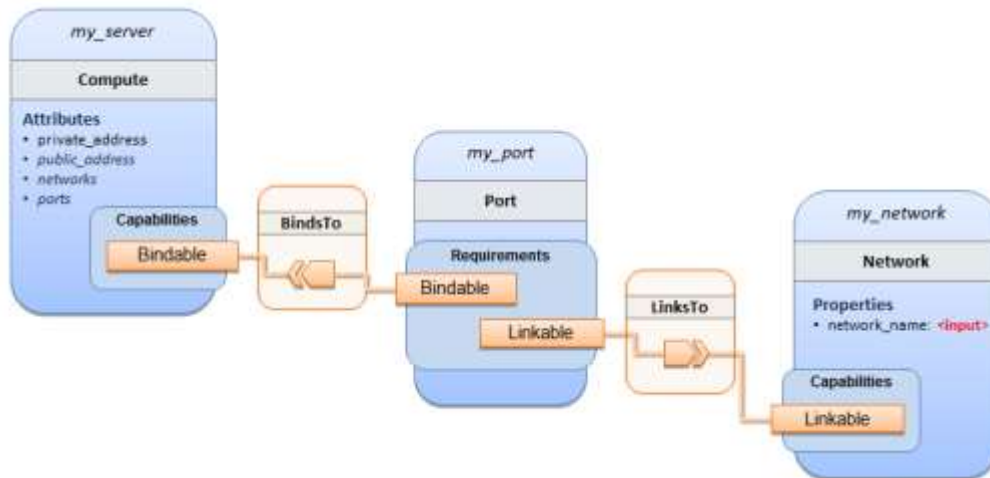
```

3938 **11.1.12 Network 2: Server bound to an existing network**

3939 **11.1.12.1 Description**

3940 This use case shows how to use a `network_name` as an input parameter to the template to allow a server
 3941 to be associated with an existing network.

3942 **11.1.12.2 Logical Diagram**



3943

3944 **11.1.12.3 Sample YAML**

```

tosca_definitions_version: toska_simple_yaml_1_0

description: >
  TOSCA simple profile with 1 server bound to an existing network

topology_template:
  inputs:
    network_name:

```

```

    type: string
    description: Network name

node_templates:
  my_server:
    type: tosca.nodes.Compute
    capabilities:
      host:
        properties:
          disk_size: 10 GB
          num_cpus: 1
          mem_size: 4096 MB
      os:
        properties:
          architecture: x86_64
          type: Linux
          distribution: CirrOS
          version: 0.3.2

  my_network:
    type: tosca.nodes.network.Network
    properties:
      network_name: { get_input: network_name }

  my_port:
    type: tosca.nodes.network.Port
    requirements:
      - binding:
          node: my_server
      - link:
          node: my_network

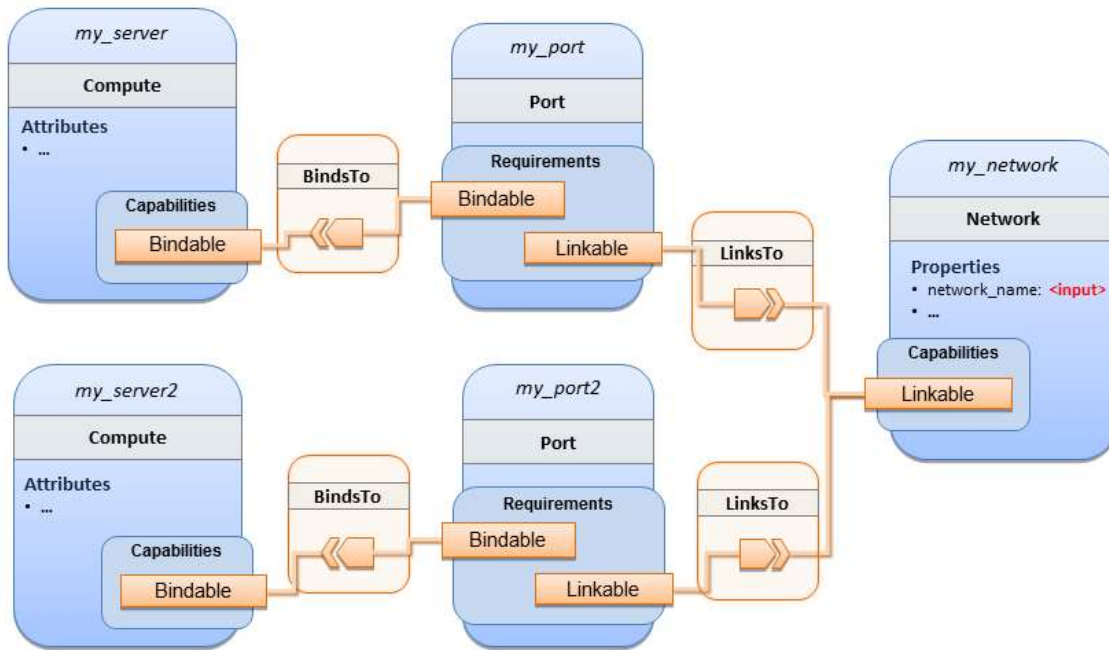
```

3945 **11.1.13 Network 3: Two servers bound to a single network**

3946 **11.1.13.1 Description**

3947 This use case shows how two servers (**Compute** nodes) can be bound to the same **Network** (node) using
 3948 two logical network **Ports**.

3949 **11.1.13.2 Logical Diagram**



3950

3951 **11.1.13.3 Sample YAML**

```

tosca_definitions_version: tosca_simple_yaml1_1_0

description: >
  TOSCA simple profile with 2 servers bound to the 1 network

topology_template:

  inputs:
    network_name:
      type: string
      description: Network name
    network_cidr:
      type: string
      default: 10.0.0.0/24
      description: CIDR for the network
    network_start_ip:
      type: string
      default: 10.0.0.100
      description: Start IP for the allocation pool
    network_end_ip:
      type: string
      default: 10.0.0.150
      description: End IP for the allocation pool

  node_templates:
    my_server:
      type: tosca.nodes.Compute
      capabilities:
        host:
          properties:

```

```

    disk_size: 10 GB
    num_cpus: 1
    mem_size: 4096 MB
  os:
    properties:
      architecture: x86_64
      type: Linux
      distribution: CirrOS
      version: 0.3.2

my_server2:
  type: toska.nodes.Compute
  capabilities:
    host:
      properties:
        disk_size: 10 GB
        num_cpus: 1
        mem_size: 4096 MB
    os:
      properties:
        architecture: x86_64
        type: Linux
        distribution: CirrOS
        version: 0.3.2

my_network:
  type: toska.nodes.network.Network
  properties:
    ip_version: 4
    cidr: { get_input: network_cidr }
    network_name: { get_input: network_name }
    start_ip: { get_input: network_start_ip }
    end_ip: { get_input: network_end_ip }

my_port:
  type: toska.nodes.network.Port
  requirements:
    - binding: my_server
    - link: my_network

my_port2:
  type: toska.nodes.network.Port
  requirements:
    - binding: my_server2
    - link: my_network

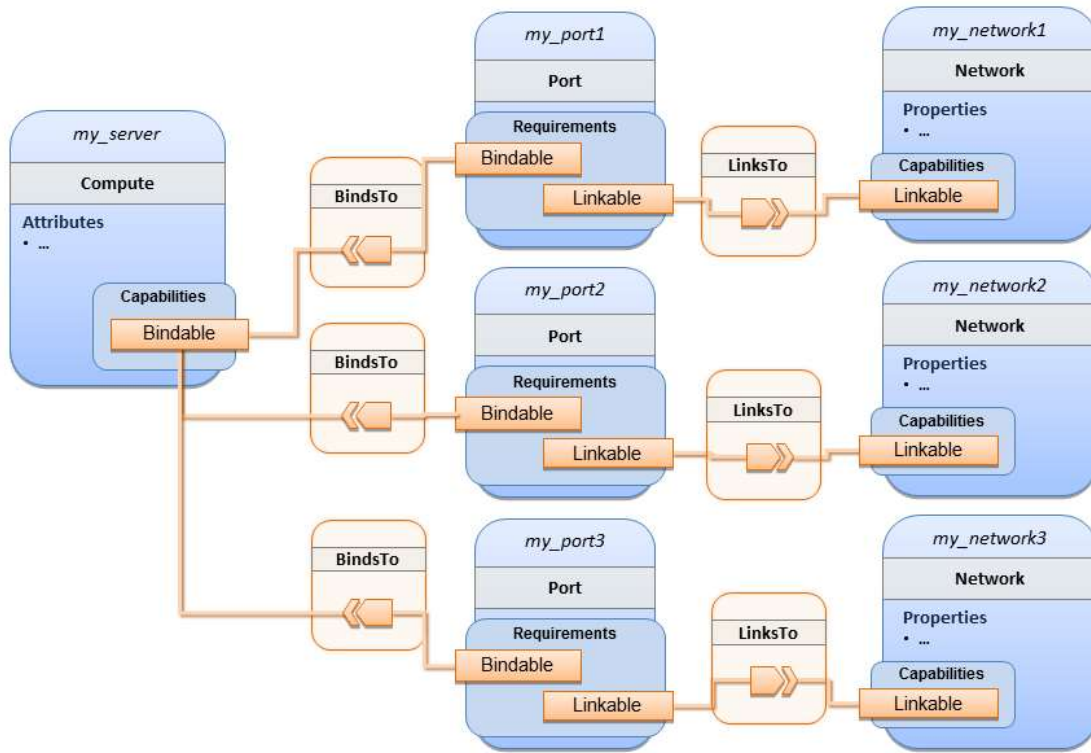
```

3952 11.1.14 Network 4: Server bound to three networks

3953 11.1.14.1 Description

3954 This use case shows how three logical networks (Network), each with its own IP address range, can be
 3955 bound to with the same server (Compute node).

3956 **11.1.14.2 Logical Diagram**



3957

3958 **11.1.14.3 Sample YAML**

```

tosca_definitions_version: tosca_simple_yaml_1_0

description: >
  TOSCA simple profile with 1 server bound to 3 networks

topology_template:

  node_templates:
    my_server:
      type: tosca.nodes.Compute
      capabilities:
        host:
          properties:
            disk_size: 10 GB
            num_cpus: 1
            mem_size: 4096 MB
        os:
          properties:
            architecture: x86_64
            type: Linux
            distribution: CirrOS
            version: 0.3.2

    my_network1:
      type: tosca.nodes.network.Network
      properties:

```

```

    cidr: '192.168.1.0/24'
    network_name: net1

my_network2:
  type: tosca.nodes.network.Network
  properties:
    cidr: '192.168.2.0/24'
    network_name: net2

my_network3:
  type: tosca.nodes.network.Network
  properties:
    cidr: '192.168.3.0/24'
    network_name: net3

my_port1:
  type: tosca.nodes.network.Port
  properties:
    order: 0
  requirements:
    - binding: my_server
    - link: my_network1

my_port2:
  type: tosca.nodes.network.Port
  properties:
    order: 1
  requirements:
    - binding: my_server
    - link: my_network2

my_port3:
  type: tosca.nodes.network.Port
  properties:
    order: 2
  requirements:
    - binding: my_server
    - link: my_network3

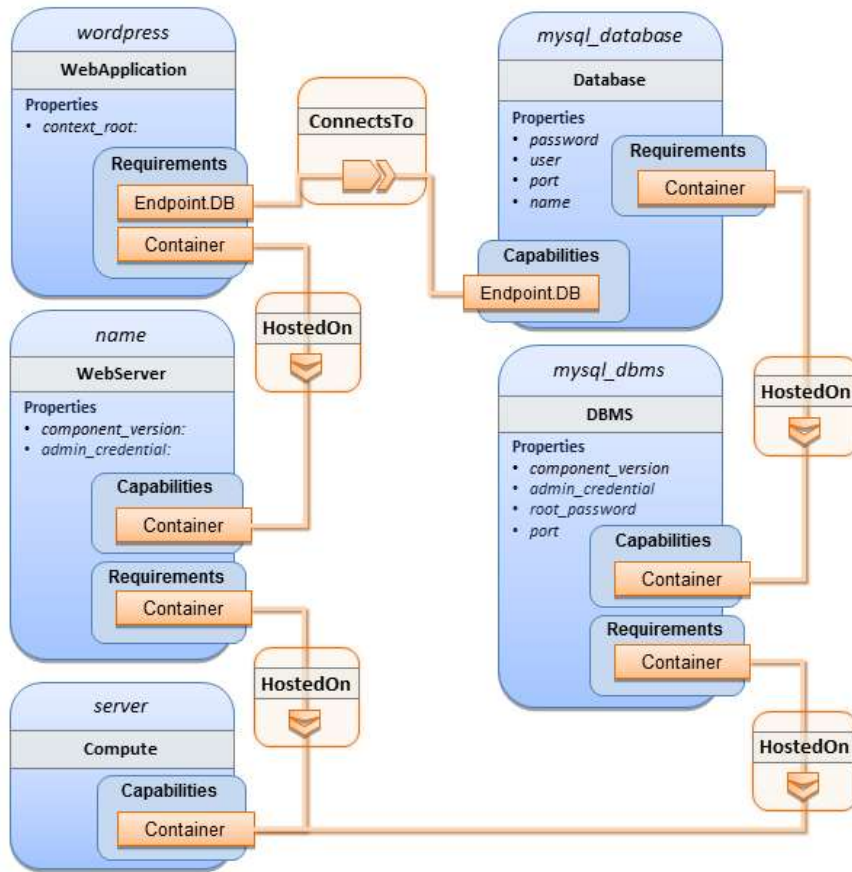
```

3959 **11.1.15 WebServer-DBMS 1: WordPress + MySQL, single instance**

3960 **11.1.15.1 Description**

3961 TOSCA simple profile service showing the WordPress web application with a MySQL database hosted on
 3962 a single server (instance).

3963 **11.1.15.2 Logical Diagram**



3964

3965 **11.1.15.3 Sample YAML**

```

tosca_definitions_version: tosca_simple_yaml1_1_0

description: >
  TOSCA simple profile with WordPress, a web server, a MySQL DBMS hosting the
  application's database content on the same server. Does not have input defaults or
  constraints.

topology_template:
  inputs:
    cpus:
      type: integer
      description: Number of CPUs for the server.
    db_name:
      type: string
      description: The name of the database.
    db_user:
      type: string
      description: The username of the DB user.
    db_pwd:
      type: string
      description: The WordPress database admin account password.
    db_root_pwd:

```

```

    type: string
    description: Root password for MySQL.
  db_port:
    type: PortDef
    description: Port for the MySQL database

node_templates:
  wordpress:
    type: tosca.nodes.WebApplication.WordPress
    properties:
      context_root: { get_input: context_root }
    requirements:
      - host: webservers
      - database_endpoint: mysql_database
    interfaces:
      Standard:
        create: wordpress_install.sh
        configure:
          implementation: wordpress_configure.sh
        inputs:
          wp_db_name: { get_property: [ mysql_database, name ] }
          wp_db_user: { get_property: [ mysql_database, user ] }
          wp_db_password: { get_property: [ mysql_database, password ] }
          # In my own template, find requirement/capability, find port property
          wp_db_port: { get_property: [ SELF, database_endpoint, port ] }

  mysql_database:
    type: Database
    properties:
      name: { get_input: db_name }
      user: { get_input: db_user }
      password: { get_input: db_pwd }
      port: { get_input: db_port }
    capabilities:
      database_endpoint:
        properties:
          port: { get_input: db_port }
    requirements:
      - host: mysql_dbms
    interfaces:
      Standard:
        configure: mysql_database_configure.sh

  mysql_dbms:
    type: DBMS
    properties:
      root_password: { get_input: db_root_pwd }
      port: { get_input: db_port }
    requirements:
      - host: server
    interfaces:
      Standard:
        inputs:
          db_root_password: { get_property: [ mysql_dbms, root_password ] }
        create: mysql_dbms_install.sh

```

```

    start: mysql_dbms_start.sh
    configure: mysql_dbms_configure.sh

webservers:
  type: WebServer
  requirements:
    - host: server
  interfaces:
    Standard:
      create: webservers_install.sh
      start: webservers_start.sh

server:
  type: Compute
  capabilities:
    host:
      properties:
        disk_size: 10 GB
        num_cpus: { get_input: cpus }
        mem_size: 4096 MB
    os:
      properties:
        architecture: x86_64
        type: linux
        distribution: fedora
        version: 17.0

outputs:
  website_url:
    description: URL for Wordpress wiki.
    value: { get_attribute: [server, public_address] }

```

3966 11.1.15.4 Sample scripts

3967 Where the referenced implementation scripts in the example above would have the following contents

3968 11.1.15.4.1 wordpress_install.sh

```
yum -y install wordpress
```

3969 11.1.15.4.2 wordpress_configure.sh

```
sed -i "/Deny from All/d" /etc/httpd/conf.d/wordpress.conf
sed -i "s/Require local/Require all granted/" /etc/httpd/conf.d/wordpress.conf
sed -i s/database_name_here/name/ /etc/wordpress/wp-config.php
sed -i s/username_here/user/ /etc/wordpress/wp-config.php
sed -i s/password_here/password/ /etc/wordpress/wp-config.php
systemctl restart httpd.service
```

3970 11.1.15.4.3 mysql_database_configure.sh

```
# Setup MySQL root password and create user
cat << EOF | mysql -u root --password=db_root_password
CREATE DATABASE name;
```

```
GRANT ALL PRIVILEGES ON name.* TO "user"@"localhost"  
IDENTIFIED BY "password";  
FLUSH PRIVILEGES;  
EXIT  
EOF
```

3971 **11.1.15.4.4 mysql_dbms_install.sh**

```
yum -y install mysql mysql-server  
# Use systemd to start MySQL server at system boot time  
systemctl enable mysqld.service
```

3972 **11.1.15.4.5 mysql_dbms_start.sh**

```
# Start the MySQL service (NOTE: may already be started at image boot time)  
systemctl start mysqld.service
```

3973 **11.1.15.4.6 mysql_dbms_configure**

```
# Set the MySQL server root password  
mysqladmin -u root password db_root_password
```

3974 **11.1.15.4.7 webserver_install.sh**

```
yum -y install httpd  
systemctl enable httpd.service
```

3975 **11.1.15.4.8 webserver_start.sh**

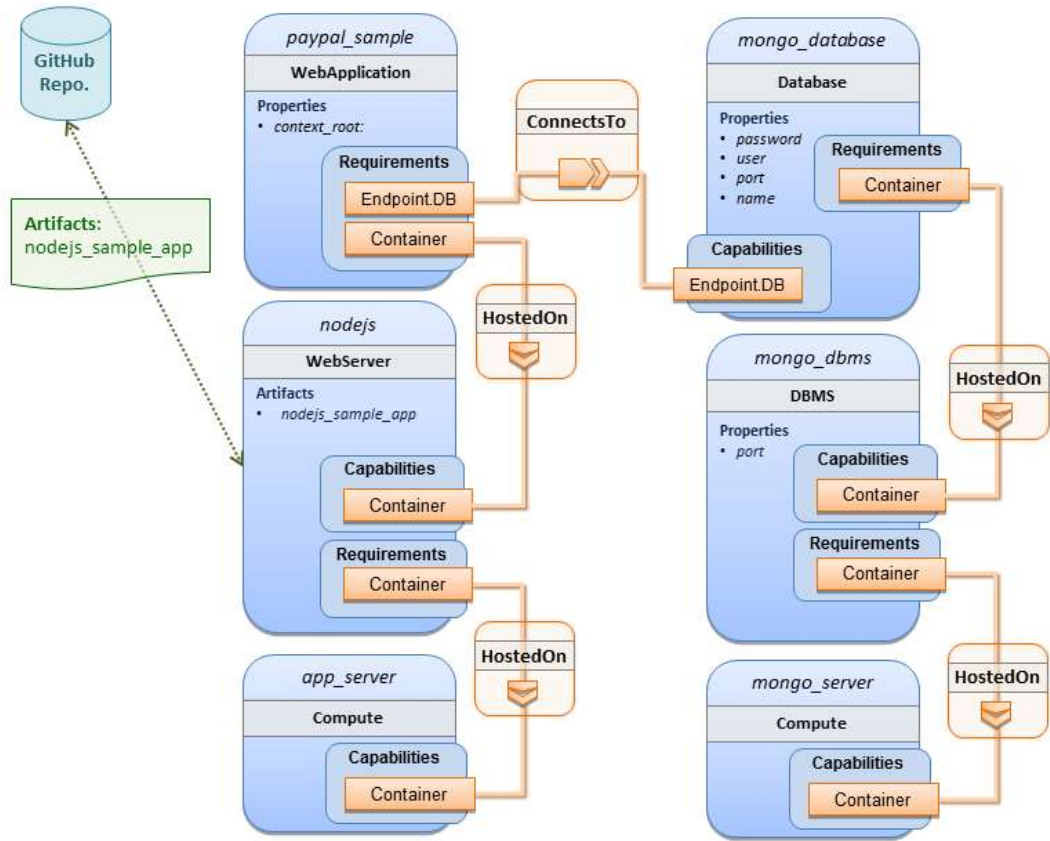
```
# Start the httpd service (NOTE: may already be started at image boot time)  
systemctl start httpd.service
```

3976 **11.1.16 WebServer-DBMS 2: Nodejs with PayPal Sample App and MongoDB** 3977 **on separate instances**

3978 **11.1.16.1 Description**

3979 This use case Instantiates a 2-tier application with Nodejs and its (PayPal sample) WebApplication on
3980 one tier which connects a MongoDB database (which stores its application data) using a ConnectsTo
3981 relationship.

3982 **11.1.16.2 Logical Diagram**



3983

3984 **11.1.16.3 Sample YAML**

```

tosca_definitions_version: tosca_simple_yaml1_1_0

description: >
  TOSCA simple profile with a nodejs web server hosting a PayPal sample application
  which connects to a mongodb database.

imports:
  - custom_types/paypalpizzastore_nodejs_app.yaml

dsl_definitions:
  ubuntu_node: &ubuntu_node
    disk_size: 10 GB
    num_cpus: { get_input: my_cpus }
    mem_size: 4096 MB
  os_capabilities: &os_capabilities
    architecture: x86_64
    type: Linux
    distribution: Ubuntu
    version: 14.04

topology_template:
  inputs:
    my_cpus:

```

```

    type: integer
    description: Number of CPUs for the server.
    constraints:
      - valid_values: [ 1, 2, 4, 8 ]
    default: 1
  github_url:
    type: string
    description: The URL to download nodejs.
    default: https://github.com/sample.git

node_templates:

  paypal_pizzastore:
    type: tosa.nodes.WebApplication.PayPalPizzaStore
    properties:
      github_url: { get_input: github_url }
    requirements:
      - host: nodejs
      - database_connection: mongo_db
    interfaces:
      Standard:
        configure:
          implementation: scripts/nodejs/configure.sh
        inputs:
          github_url: { get_property: [ SELF, github_url ] }
          mongodb_ip: { get_attribute: [ mongo_server, private_address ] }
        start: scriptsscripts/nodejs/start.sh

  nodejs:
    type: tosa.nodes.WebServer.Nodejs
    requirements:
      - host: app_server
    interfaces:
      Standard:
        create: scripts/nodejs/create.sh

  mongo_db:
    type: tosa.nodes.Database
    requirements:
      - host: mongo_dbms
    interfaces:
      Standard:
        create: create_database.sh

  mongo_dbms:
    type: tosa.nodes.DBMS
    requirements:
      - host: mongo_server
    properties:
      port: 27017
    interfaces:
      tosa.interfaces.node.lifecycle.Standard:
        create: mongodb/create.sh
        configure:
          implementation: mongodb/config.sh

```

```

    inputs:
      mongodb_ip: { get_attribute: [mongo_server, private_address] }
      start: mongodb/start.sh

  mongo_server:
    type: tosca.nodes.Compute
    capabilities:
      os:
        properties: *os_capabilities
      host:
        properties: *ubuntu_node

  app_server:
    type: tosca.nodes.Compute
    capabilities:
      os:
        properties: *os_capabilities
      host:
        properties: *ubuntu_node

  outputs:
    nodejs_url:
      description: URL for the nodejs server, http://<IP>:3000
      value: { get_attribute: [app_server, private_address] }
    mongodb_url:
      description: URL for the mongodb server.
      value: { get_attribute: [ mongo_server, private_address ] }

```

3985 **11.1.16.4 Notes:**

- 3986
- Scripts referenced in this example are assumed to be placed by the TOSCA orchestrator in the relative directory declared in TOSCA.meta of the TOSCA CSAR file.
- 3987

3988 **11.1.17 Multi-Tier-1: Elasticsearch, Logstash, Kibana (ELK) use case with**

3989 **multiple instances**

3990 **11.1.17.1 Description**

3991 TOSCA simple profile service showing the Nodejs, MongoDB, Elasticsearch, Logstash, Kibana, rsyslog

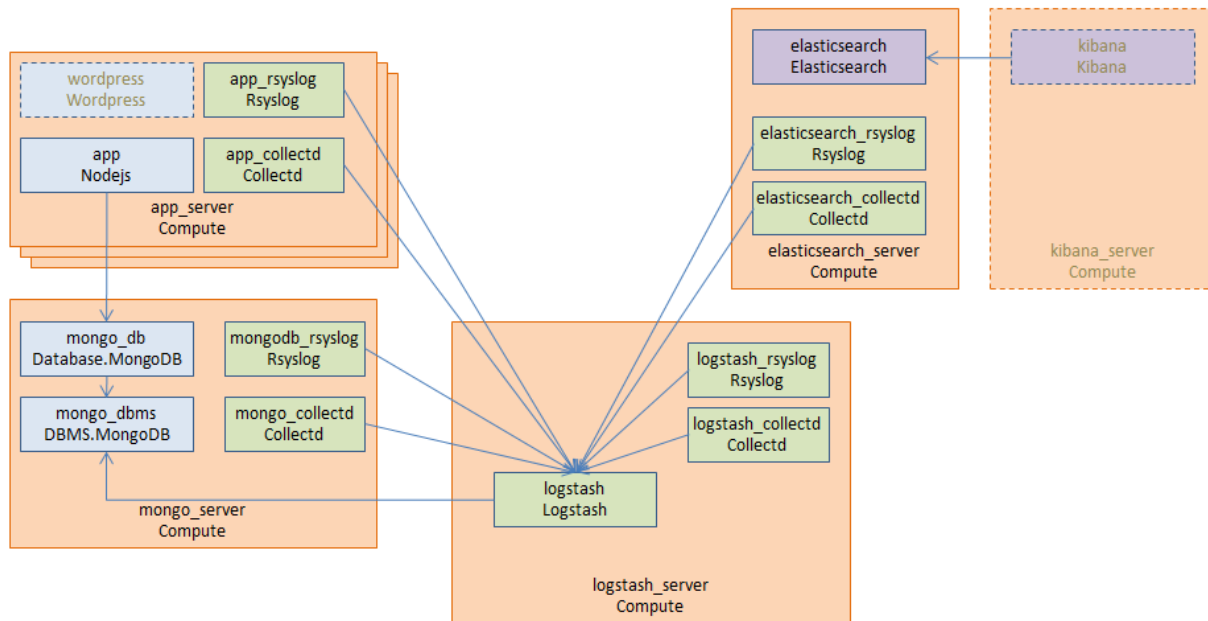
3992 and collectd installed on a different server (instance).

3993

3994 This use case also demonstrates:

- Use of TOSCA macros or dsl_definitions
 - Multiple **SoftwareComponents** hosted on same Compute node
 - Multiple tiers communicating to each other over ConnectsTo using Configure interface.
- 3995
- 3996
- 3997

3998 **11.1.17.2 Logical Diagram**



3999

4000 **11.1.17.3 Sample YAML**

4001 **11.1.17.3.1 Master Service Template application (Entry-Definitions)**

4002 TheThe following YAML is the primary template (i.e., the Entry-Definition) for the overall use case. The
 4003 imported YAML for the various subcomponents are not shown here for brevity.

4004

```
tosca_definitions_version: tosca_simple_yaml_1_0

description: >
  This TOSCA simple profile deploys nodejs, mongodb, elasticsearch, logstash and kibana each on a separate server with monitoring enabled for nodejs server where a sample nodejs application is running. The syslog and collectd are installed on a nodejs server.

imports:
- paypalpizzastore_nodejs_app.yaml
- elasticsearch.yaml
- logstash.yaml
- kibana.yaml
- collectd.yaml
- rsyslog.yaml

dsl_definitions:
  host_capabilities: &host_capabilities
    # container properties (flavor)
    disk_size: 10 GB
    num_cpus: { get_input: my_cpus }
    mem_size: 4096 MB
    os_capabilities: &os_capabilities
```



```

architecture: x86_64
type: Linux
distribution: Ubuntu
version: 14.04

topology_template:
  inputs:
    my_cpus:
      type: integer
      description: Number of CPUs for the server.
      constraints:
        - valid_values: [ 1, 2, 4, 8 ]
    github_url:
      type: string
      description: The URL to download nodejs.
      default: https://github.com/sample.git

  node_templates:
    paypal_pizzastore:
      type: toska.nodes.WebApplication.PayPalPizzaStore
      properties:
        github_url: { get_input: github_url }
      requirements:
        - host: nodejs
        - database_connection: mongo_db
      interfaces:
        Standard:
          configure:
            implementation: scripts/nodejs/configure.sh
            inputs:
              github_url: { get_property: [ SELF, github_url ] }
              mongodb_ip: { get_attribute: [mongo_server, private_address] }
            start: scripts/nodejs/start.sh

    nodejs:
      type: toska.nodes.WebServer.Nodejs
      requirements:
        - host: app_server
      interfaces:
        Standard:
          create: scripts/nodejs/create.sh

    mongo_db:
      type: toska.nodes.Database
      requirements:
        - host: mongo_dbms
      interfaces:
        Standard:
          create: create_database.sh

    mongo_dbms:
      type: toska.nodes.DBMS
      requirements:
        - host: mongo_server
      interfaces:

```

```

tosca.interfaces.node.lifecycle.Standard:
  create: scripts/mongodb/create.sh
  configure:
    implementation: scripts/mongodb/config.sh
  inputs:
    mongodb_ip: { get_attribute: [mongo_server, ip_address] }
  start: scripts/mongodb/start.sh

elasticsearch:
  type: tosca.nodes.SoftwareComponent.Elasticsearch
  requirements:
    - host: elasticsearch_server
  interfaces:
    tosca.interfaces.node.lifecycle.Standard:
      create: scripts/elasticsearch/create.sh
      start: scripts/elasticsearch/start.sh
logstash:
  type: tosca.nodes.SoftwareComponent.Logstash
  requirements:
    - host: logstash_server
    - search_endpoint: elasticsearch
  interfaces:
    tosca.interfaces.relationship.Configure:
      pre_configure_source:
        implementation: python/logstash/configure_elasticsearch.py
      input:
        elasticsearch_ip: { get_attribute: [elasticsearch_server,
ip_address] }
    interfaces:
      tosca.interfaces.node.lifecycle.Standard:
        create: scripts/logstash/create.sh
        configure: scripts/logstash/config.sh
        start: scripts/logstash/start.sh

kibana:
  type: tosca.nodes.SoftwareComponent.Kibana
  requirements:
    - host: kibana_server
    - search_endpoint: elasticsearch
  interfaces:
    tosca.interfaces.node.lifecycle.Standard:
      create: scripts/kibana/create.sh
      configure:
        implementation: scripts/kibana/config.sh
      input:
        elasticsearch_ip: { get_attribute: [elasticsearch_server, ip_address]
}
        kibana_ip: { get_attribute: [kibana_server, ip_address] }
      start: scripts/kibana/start.sh

app_collectd:
  type: tosca.nodes.SoftwareComponent.Collectd
  requirements:
    - host: app_server
    - collectd_endpoint: logstash

```

```

    interfaces:
      tosca.interfaces.relationship.Configure:
        pre_configure_target:
          implementation: python/logstash/configure_collectd.py
interfaces:
  tosca.interfaces.node.lifecycle.Standard:
    create: scripts/collectd/create.sh
    configure:
      implementation: python/collectd/config.py
      input:
        logstash_ip: { get_attribute: [logstash_server, ip_address] }
    start: scripts/collectd/start.sh

app_rsyslog:
  type: tosca.nodes.SoftwareComponent.Rsyslog
  requirements:
    - host: app_server
    - rsyslog_endpoint: logstash
  interfaces:
    tosca.interfaces.relationship.Configure:
      pre_configure_target:
        implementation: python/logstash/configure_rsyslog.py
  interfaces:
    tosca.interfaces.node.lifecycle.Standard:
      create: scripts/rsyslog/create.sh
      configure:
        implementation: scripts/rsyslog/config.sh
      input:
        logstash_ip: { get_attribute: [logstash_server, ip_address] }
      start: scripts/rsyslog/start.sh

app_server:
  type: tosca.nodes.Compute
  capabilities:
    host:
      properties: *host_capabilities
    os:
      properties: *os_capabilities

mongo_server:
  type: tosca.nodes.Compute
  capabilities:
    host:
      properties: *host_capabilities
    os:
      properties: *os_capabilities

elasticsearch_server:
  type: tosca.nodes.Compute
  capabilities:
    host:
      properties: *host_capabilities
    os:
      properties: *os_capabilities

```

```

logstash_server:
  type: toska.nodes.Compute
  capabilities:
    host:
      properties: *host_capabilities
    os:
      properties: *os_capabilities

kibana_server:
  type: toska.nodes.Compute
  capabilities:
    host:
      properties: *host_capabilities
    os:
      properties: *os_capabilities

outputs:
  nodejs_url:
    description: URL for the nodejs server.
    value: { get_attribute: [ app_server, private_address ] }
  mongodb_url:
    description: URL for the mongodb server.
    value: { get_attribute: [ mongo_server, private_address ] }
  elasticsearch_url:
    description: URL for the elasticsearch server.
    value: { get_attribute: [ elasticsearch_server, private_address ] }
  logstash_url:
    description: URL for the logstash server.
    value: { get_attribute: [ logstash_server, private_address ] }
  kibana_url:
    description: URL for the kibana server.
    value: { get_attribute: [ kibana_server, private_address ] }

```

4005 11.1.17.4 Sample scripts

4006 Where the referenced implementation scripts in the example above would have the following contents

4007 11.1.18 Container-1: Containers using Docker single Compute instance 4008 (Containers only)

4009 11.1.18.1 Description

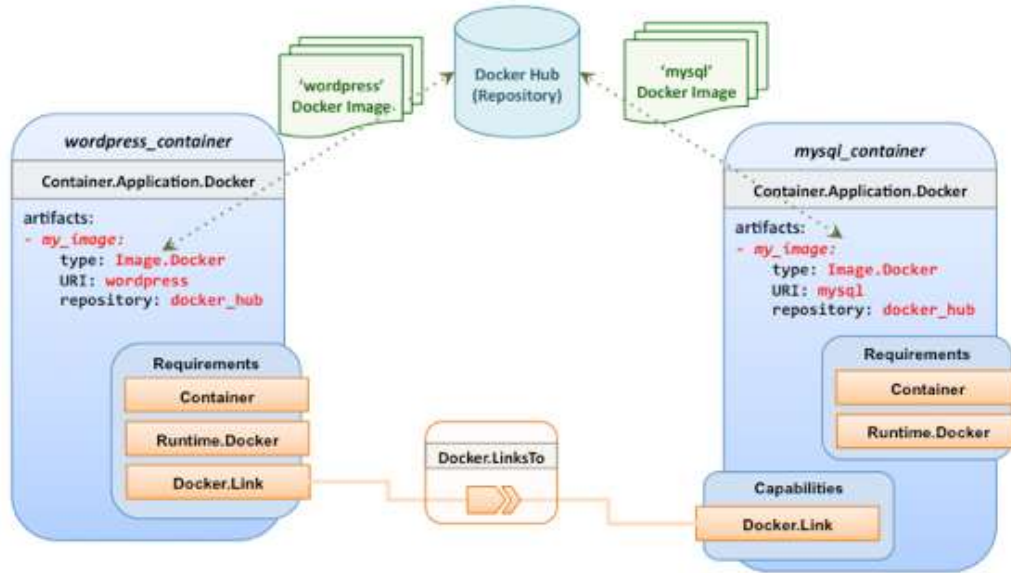
4010 This use case shows a minimal description of two Container nodes (only) providing their Docker
4011 Requirements allowing platform (orchestrator) to select/provide the underlying Docker implementation
4012 (Capability). Specifically, wordpress and mysql Docker images are referenced from Docker Hub.

4013

4014 This use case also demonstrates:

- 4015 • Abstract description of Requirements (i.e., Container and Docker) allowing platform to
- 4016 dynamically select the appropriate runtime Capabilities that match.
- 4017 • Use of external repository (Docker Hub) to reference image artifact.

4018 **11.1.18.2 Logical Diagram**



4019

4020 **11.1.18.3 Sample YAML**

4021 **11.1.18.3.1 Two Docker “Container” nodes (Only) with Docker Requirements**

```

tosca_definitions_version: tosca_simple_yaml_1_0

description: >
  TOSCA simple profile with wordpress, web server and mysql on the same server.

# Repositories to retrieve code artifacts from
repositories:
  docker_hub: https://registry.hub.docker.com/

topology_template:

  inputs:
    wp_host_port:
      type: integer
      description: The host port that maps to port 80 of the WordPress container.
    db_root_pwd:
      type: string
      description: Root password for MySQL.

  node_templates:
    # The MYSQL container based on official MySQL image in Docker hub
    mysql_container:
      type: tosca.nodes.Container.Application.Docker
      capabilities:
        # This is a capability that would mimic the Docker -link feature
        database_link: tosca.capabilities.Docker.Link
      artifacts:
        my_image:
          file: mysql
  
```

```

    type: tosca.artifacts.Deployment.Image.Container.Docker
    repository: docker_hub
  interfaces:
    Standard:
      create:
        implementation: my_image
        inputs:
          db_root_password: { get_input: db_root_pwd }

# The WordPress container based on official WordPress image in Docker hub
wordpress_container:
  type: tosca.nodes.Container.Application.Docker
  requirements:
    - database_link: mysql_container
  artifacts:
    my_image:
      file: wordpress
      type: tosca.artifacts.Deployment.Image.Container.Docker
      repository: docker_hub
  interfaces:
    Standard:
      create:
        implementation: my_image
        inputs:
          host_port: { get_input: wp_host_port }

```

4022

4023 12 TOSCA Policies

4024 This section is **non-normative** and describes the approach TOSCA Simple Profile plans to take for policy
4025 description with TOSCA Service Templates. In addition, it explores how existing TOSCA Policy Types
4026 and definitions might be applied in the future to express operational policy use cases.

4027 12.1 A declarative approach

4028 TOSCA Policies are a type of requirement that govern use or access to resources which can be
4029 expressed independently from specific applications (or their resources) and whose fulfillment is not
4030 discretely expressed in the application's topology (i.e., via TOSCA Capabilities).

4031
4032 TOSCA deems it not desirable for a declarative model to encourage external intervention for resolving
4033 policy issues (i.e., via imperative mechanisms external to the Cloud). Instead, the Cloud provider is
4034 deemed to be in the best position to detect when policy conditions are triggered, analyze the affected
4035 resources and enforce the policy against the allowable actions declared within the policy itself.

4036 12.1.1 Declarative considerations

- 4037 • Natural language rules are not realistic, too much to represent in our specification; however, regular
4038 expressions can be used that include simple operations and operands that include symbolic names
4039 for TOSCA metamodel entities, properties and attributes.
- 4040 • Complex rules can actually be directed to an external policy engine (to check for violation) returns
4041 true/false then policy says what to do (trigger or action).
- 4042 • Actions/Triggers could be:
 - 4043 • Autonomic/Platform corrects against user-supplied criteria
 - 4044 • External monitoring service could be utilized to monitor policy rules/conditions against metrics,
4045 the monitoring service could coordinate corrective actions with external services (perhaps
4046 Workflow engines that can analyze the application and interact with the TOSCA instance model).

4047 12.2 Consideration of Event, Condition and Action

4048 12.3 Types of policies

4049 Policies typically address two major areas of concern for customer workloads:

- 4050 • **Access Control** – assures user and service access to controlled resources are governed by
4051 rules which determine general access permission (i.e., allow or deny) and conditional access
4052 dependent on other considerations (e.g., organization role, time of day, geographic location, etc.).
- 4053 • **Placement** – assures affinity (or anti-affinity) of deployed applications and their resources; that is,
4054 what is allowed to be placed where within a Cloud provider's infrastructure.
- 4055 • **Quality-of-Service** (and continuity) - assures performance of software components (perhaps
4056 captured as quantifiable, measure components within an SLA) along with consideration for
4057 scaling and failover.

4058 12.3.1 Access control policies

4059 Although TOSCA Policy definitions could be used to express and convey access control policies,
4060 definitions of policies in this area are out of scope for this specification. At this time, TOSCA encourages
4061 organizations that already have standards that express policy for access control to provide their own
4062 guidance on how to use their standard with TOSCA.

4063 12.3.2 Placement policies

4064 There must be control mechanisms in place that can be part of these patterns that accept governance
4065 policies that allow control expressions of what is allowed when placing, scaling and managing the
4066 applications that are enforceable and verifiable in Cloud.

4067

4068 These policies need to consider the following:

- 4069 • Regulated industries need applications to control placement (deployment) of applications to
4070 different countries or regions (i.e., different logical geographical boundaries).

4071 12.3.2.1 Placement for governance concerns

4072 In general, companies and individuals have security concerns along with general “loss of control” issues
4073 when considering deploying and hosting their highly valued application and data to the Cloud. They want
4074 to control placement perhaps to ensure their applications are only placed in datacenter they trust or
4075 assure that their applications and data are not placed on shared resources (i.e., not co-tenanted).

4076

4077 In addition, companies that are related to highly regulated industries where compliance with government,
4078 industry and corporate policies is paramount. In these cases, having the ability to control placement of
4079 applications is an especially significant consideration and a prerequisite for automated orchestration.

4080 12.3.2.2 Placement for failover

4081 Companies realize that their day-to-day business must continue on through unforeseen disasters that
4082 might disable instances of the applications and data at or on specific data centers, networks or servers.
4083 They need to be able to convey placement policies for their software applications and data that mitigate
4084 risk of disaster by assuring these cloud assets are deployed strategically in different physical locations.
4085 Such policies need to consider placement across geographic locations as wide as countries, regions,
4086 datacenters, as well as granular placement on a network, server or device within the same physical
4087 datacenter. Cloud providers must be able to not only enforce these policies but provide robust and
4088 seamless failover such that a disaster’s impact is never perceived by the end user.

4089 12.3.3 Quality-of-Service (QoS) policies

4090 Quality-of-Service (apart from failover placement considerations) typically assures that software
4091 applications and data are available and performant to the end users. This is usually something that is
4092 measurable in terms of end-user responsiveness (or response time) and often qualified in SLAs
4093 established between the Cloud provider and customer. These QoS aspects can be taken from SLAs and
4094 legal agreements and further encoded as performance policies associated with the actual applications
4095 and data when they are deployed. It is assumed that Cloud provider is able to detect high utilization (or
4096 usage load) on these applications and data that deviate from these performance policies and is able to
4097 bring them back into compliance.

4098

4099 12.4 Policy relationship considerations

- 4100 • Performance policies can be related to scalability policies. Scalability policies tell the Cloud provider
4101 exactly **how** to scale applications and data when they detect an application’s performance policy is
4102 (or about to be) violated (or triggered).
- 4103 • Scalability policies in turn are related to placement policies which govern **where** the application and
4104 data can be scaled to.
- 4105 • There are general “tenant” considerations that restrict what resources are available to applications
4106 and data based upon the contract a customer has with the Cloud provider. This includes other

4107 constraints imposed by legal agreements or SLAs that are not encoded programmatically or
4108 associated directly with actual application or data..

4109 **12.5 Use Cases**

4110 This section includes some initial operation policy use cases that we wish to describe using the TOSCA
4111 metamodel. More policy work will be done in future versions of the TOSCA Simple Profile in YAML
4112 specification.

4113 **12.5.1 Placement**

4114 **12.5.1.1 Use Case 1: Simple placement for failover**

4115 **12.5.1.1.1 Description**

4116 This use case shows a failover policy to keep at least 3 copies running in separate containers. In this
4117 simple case, the specific containers to use (or name is not important; the Cloud provider must assure
4118 placement separation (anti-affinity) in three physically separate containers.

4119 **12.5.1.1.2 Features**

4120 This use case introduces the following policy features:

- 4121 • Simple separation on different “compute” nodes (up to discretion of provider).
- 4122 • Simple separation by region (a logical container type) using an allowed list of region names
4123 relative to the provider.
 - 4124 ○ Also, shows that set of allowed “regions” (containers) can be greater than the number of
4125 containers requested.

4126 **12.5.1.1.3 Logical Diagram**

4127 Sample YAML: Compute separation

```
failover_policy_1:  
  type: tosca.policy.placement.Antilocate  
  description: My placement policy for Compute node separation  
  properties:  
    # 3 diff target containers  
    container_type: Compute  
    container_number: 3
```

4128 **12.5.1.1.4 Notes**

- 4129 • There may be availability (constraints) considerations especially if these policies are applied to
4130 “clusters”.
- 4131 • There may be future considerations for controlling max # of instances per container.

4132 **12.5.1.2 Use Case 2: Controlled placement by region**

4133 **12.5.1.2.1 Description**

4134 This use case demonstrates the use of named “containers” which could represent the following:

- 4135 • Datacenter regions
- 4136 • Geographic regions (e.g., cities, municipalities, states, countries, etc.)
- 4137 • Commercial regions (e.g., North America, Eastern Europe, Asia Pacific, etc.)

4138 **12.5.1.2.2 Features**

4139 This use case introduces the following policy features:

- 4140
 - Separation of resources (i.e., TOSCA nodes) by logical regions, or zones.

4141 **12.5.1.2.3 Sample YAML: Region separation amongst named set of regions**

```
failover_policy_2:
  type: toska.policy.placement
  description: My failover policy with allowed target regions (logical containers)
  properties:
    container type: region
    container_number: 3
    # If "containers" keyname is provided, they represent the allowed set
    # of target containers to use for placement for .
    containers: [ region1, region2, region3, region4 ]
```

4142 **12.5.1.3 Use Case 3: Co-locate based upon Compute affinity**

4143 **12.5.1.3.1 Description**

4144 Nodes that need to be co-located to achieve optimal performance based upon access to similar
4145 Infrastructure (IaaS) resource types (i.e., Compute, Network and/or Storage).

4146

4147 This use case demonstrates the co-location based upon Compute resource affinity; however, the same
4148 approach could be taken for Network as or Storage affinity as well. :

4149 **12.5.1.3.2 Features**

4150 This use case introduces the following policy features:

- 4151
 - Node placement based upon Compute resource affinity.

4152 **12.5.1.4 Notes**

- 4153
 - The concept of placement based upon IaaS resource utilization is not future-thinking, as Cloud
4154 should guarantee equivalent performance of application performance regardless of placement.
4155 That is, all network access between application nodes and underlying Compute or Storage should
4156 have equivalent performance (e.g., network bandwidth, network or storage access time, CPU
4157 speed, etc.).

4158 **12.5.1.4.1 Sample YAML: Region separation amongst named set of regions**

```
keep_together_policy:
  type: toska.policy.placement.Colocate
  description: Keep associated nodes (groups of nodes) based upon Compute
  properties:
    affinity: Compute
```

4159 12.5.2 Scaling

4160 12.5.2.1 Use Case 1: Simple node autoscale

4161 12.5.2.1.1 Description

4162 Start with X nodes and scale up to Y nodes, capability to do this from a dashboard for example.

4163 12.5.2.1.2 Features

4164 This use case introduces the following policy features:

- 4165 • Basic autoscaling policy

4166 12.5.2.1.3 Sample YAML

```
my_scaling_policy_1:
  type: toska.policy.scaling
  description: Simple node autoscaling
  properties:
    min_instances: <integer>
    max_instances: <integer>
    default_instances: <integer>
    increment: <integer>
```

4167 12.5.2.1.4 Notes

- 4168 • Assume horizontal scaling for this use case
 - 4169 ○ Horizontal scaling, implies “stack-level” control using Compute nodes to define a “stack”
 - 4170 (i.e., The Compute node's entire HostedOn relationship dependency graph is considered
 - 4171 part of its “stack”)
- 4172 • Assume Compute node has a SoftwareComponent that represents a VM application.
- 4173 • Availability Zones (and Regions if not same) need to be considered in further
- 4174 use cases.
- 4175 • If metrics are introduced, there is a control-loop (that monitors). Autoscaling is a special concept
- 4176 that includes these considerations.
- 4177 • Mixed placement and scaling use cases need to be considered:
 - 4178 ○ *Example:* Compute1 and Compute2 are 2 node templates. Compute1 has 10 instances, 5
 - 4179 in one region 5 in other region.

4180

13 Conformance

4181

13.1 Conformance Targets

4182

The implementations subject to conformance are those introduced in Section 11.3 “Implementations”.

4183

They are listed here for convenience:

4184

- TOSCA YAML service template

4185

- TOSCA processor

4186

- TOSCA orchestrator (also called orchestration engine)

4187

- TOSCA generator

4188

- TOSCA archive

4189

13.2 Conformance Clause 1: TOSCA YAML service template

4190

A document conforms to this specification as TOSCA YAML service template if it satisfies all the

4191

statements below:

4192

- (a) It is valid according to the grammar, rules and requirements defined in section 3 “TOSCA Simple Profile definitions in YAML”.

4193

4194

- (b) When using functions defined in section 4 “TOSCA functions”, it is valid according to the grammar specified for these functions.

4195

4196

- (c) When using or referring to data types, artifact types, capability types, interface types, node types, relationship types, group types, policy types defined in section 5 “TOSCA normative type definitions”, it is valid according to the definitions given in section 5.

4197

4198

4199

13.3 Conformance Clause 2: TOSCA processor

4200

A processor or program conforms to this specification as TOSCA processor if it satisfies all the

4201

statements below:

4202

- (a) It can parse and recognize the elements of any conforming TOSCA YAML service template, and generates errors for those documents that fail to conform as TOSCA YAML service template while clearly intending to.

4203

4204

4205

- (b) It implements the requirements and semantics associated with the definitions and grammar in section 3 “TOSCA Simple Profile definitions in YAML”, including those listed in the “additional requirements” subsections.

4206

4207

4208

- (c) It resolves the imports, either explicit or implicit, as described in section 3 “TOSCA Simple Profile definitions in YAML”.

4209

4210

- (d) It generates errors as required in error cases described in sections 3.1 (TOSCA Namespace URI and alias), 3.2 (Parameter and property type) and 3.6 (Type-specific definitions).

4211

4212

- (e) It normalizes string values as described in section 5.4.9.3 (Additional Requirements)

4213

13.4 Conformance Clause 3: TOSCA orchestrator

4214

A processor or program conforms to this specification as TOSCA orchestrator if it satisfies all the

4215

statements below:

4216

- (a) It is conforming as a TOSCA Processor as defined in conformance clause 2: TOSCA Processor.

4217

- (b) It can process all types of artifact described in section 5.3 “Artifact types” according to the rules and grammars in this section.

4218

- 4219 (c) It can process TOSCA archives as intended in section 6 “TOSCA Cloud Service Archive (CSAR)
4220 format” and other related normative sections.
4221 (d) It can understand and process the functions defined in section 4 “TOSCA functions” according to
4222 their rules and semantics.
4223 (e) It can understand and process the normative type definitions according to their semantics and
4224 requirements as described in section 5 “TOSCA normative type definitions”.
4225 (f) It can understand and process the networking types and semantics defined in section 7 “TOSCA
4226 Networking”.
4227 (g) It generates errors as required in error cases described in sections 2.10 (Using node template
4228 substitution for chaining subsystems), 5.4 (Capabilities Types) and 5.7 (Interface Types).

4229 **13.5 Conformance Clause 4: TOSCA generator**

4230 A processor or program conforms to this specification as TOSCA generator if it satisfies at least one of
4231 the statements below:

- 4232 (a) When requested to generate a TOSCA service template, it always produces a conforming
4233 TOSCA service template, as defined in Clause 1: TOSCA YAML service template,
4234 (b) When requested to generate a TOSCA archive, it always produces a conforming TOSCA archive,
4235 as defined in Clause 5: TOSCA archive.

4236 **13.6 Conformance Clause 5: TOSCA archive**

4237 A package artifact conforms to this specification as TOSCA archive if it satisfies all the statements below:

- 4238 (a) It is valid according to the structure and rules defined in section 6 “TOSCA Cloud Service Archive
4239 (CSAR) format”.

4240

Appendix A. Known Extensions to TOSCA v1.0

4241 The following items will need to be reflected in the TOSCA (XML) specification to allow for isomorphic
4242 mapping between the XML and YAML service templates.

4243 A.1 Model Changes

- 4244 • The “TOSCA Simple ‘Hello World’” example introduces this concept in Section 2. Specifically, a VM
4245 image assumed to accessible by the cloud provider.
- 4246 • Introduce template Input and Output parameters
- 4247 • The “Template with input and output parameter” example introduces concept in Section 2.1.1.
 - 4248 • “Inputs” could be mapped to BoundaryDefinitions in TOSCA v1.0. Maybe needs some usability
4249 enhancement and better description.
 - 4250 • “outputs” are a new feature.
- 4251 • Grouping of Node Templates
 - 4252 • This was part of original TOSCA proposal, but removed early on from v1.0 This allows grouping
4253 of node templates that have some type of logically managed together as a group (perhaps to
4254 apply a scaling or placement policy).
- 4255 • Lifecycle Operation definition independent/separate from Node Types or Relationship types (allows
4256 reuse). For now, we added definitions for “node.lifecycle” and “relationship.lifecycle”.
- 4257 • Override of Interfaces (operations) in the Node Template.
- 4258 • Service Template Naming/Versioning
 - 4259 • Should include TOSCA spec. (or profile) version number (as part of namespace)
- 4260 • Allow the referencing artifacts using a URL (e.g., as a property value).
- 4261 • Repository definitions in Service Template.
- 4262 • Substitution mappings for Topology template.
- 4263 • Addition of Group Type, Policy Type, Group def., Policy def. along with normative TOSCA base types
4264 for policies and groups.

4265 A.2 Normative Types

- 4266 • Constraints
 - 4267 • constraint clauses, regex
- 4268 • Types / Property / Parameters
 - 4269 • list, map, range, scalar-unit types
 - 4270 • Includes YAML intrinsic types
 - 4271 • NetworkInfo, PortInfo, PortDef, PortSpec, Credential
 - 4272 • TOSCA Version based on Maven
- 4273 • Node
 - 4274 • Root, Compute, ObjectStorage, BlockStorage, Network, Port, SoftwareComponent,
4275 WebServer, WebApplicaton, DBMS, Database, Container, and others
- 4276 • Relationship
 - 4277 • Root, DependsOn, HostedOn, ConnectsTo, AttachesTo, RoutesTo, BindsTo, LinksTo and
4278 others
- 4279 • Artifact
 - 4280 • Deployment: Image Types (e.g., VM, Container), ZIP, TAR, etc.
 - 4281 • Implementation: File, Bash, Python, etc.
- 4282 • Requirements

- 4283
 - None
- 4284
 - Capabilities
- 4285
 - Container, Endpoint, Attachment, Scalable, ...
- 4286
 - Lifecycle
- 4287
 - Standard (for Node Types)
- 4288
 - Configure (for Relationship Types)
- 4289
 - Functions
- 4290
 - get_input, get_attribute, get_property, get_nodes_of_type, get_operation_output and others
- 4291
 - concat, token
- 4292
 - get_artifact
- 4293
 - Groups
- 4294
 - Root
- 4295
 - Policies
- 4296
 - Root, Placement, Scaling, Update, Performance
- 4297
 - Workflow

4298

Appendix B. Acknowledgments

4299 The following individuals have participated in the creation of this specification and are gratefully
4300 acknowledged:

4301 **Contributors:**

4302 Avi Vachnis (avi.vachnis@alcatel-lucent.com), Alcatel-Lucent

4303 Chris Lauwers (lauwers@ubicity.com)

4304 Derek Palma (dpalma@vnomnic.com), Vnomic

4305 Frank Leymann (Frank.Leymann@informatik.uni-stuttgart.de), Univ. of Stuttgart

4306 Gerd Breiter (gbreiter@de.ibm.com), IBM

4307 Hemal Surti (hsurti@cisco.com), Cisco

4308 Ifat Afek (ifat.afek@alcatel-lucent.com), Alcatel-Lucent

4309 Idan Moyal, (idan@gigaspaces.com), Gigaspaces

4310 Jacques Durand (jdurand@us.fujitsu.com), Fujitsu

4311 Jin Qin, (chin.qinjin@huawei.com), Huawei

4312 Jeremy Hess, (jeremy@gigaspaces.com) , Gigaspaces

4313 John Crandall, (mailto:jcrandal@brocade.com), Brocade

4314 Juergen Meynert (juergen.meynert@ts.fujitsu.com), Fujitsu

4315 Kapil Thangavelu (kapil.thangavelu@canonical.com), Canonical

4316 Karsten Beins (karsten.beins@ts.fujitsu.com), Fujitsu

4317 Kevin Wilson (kevin.l.wilson@hp.com), HP

4318 Krishna Raman (kraman@redhat.com), Red Hat

4319 Luc Boutier (luc.boutier@fastconnect.fr), FastConnect

4320 Luca Gioppo, (luca.gioppo@csi.it), CSI-Piemonte

4321 Matt Rutkowski (mrutkows@us.ibm.com), IBM

4322 Moshe Elisha (moshe.elisha@alcatel-lucent.com), Alcatel-Lucent

4323 Nate Finch (nate.finch@canonical.com), Canonical

4324 Nikunj Nemani (nnemani@vmware.com), VMware

4325 Richard Probst (richard.probst@sap.com), SAP AG

4326 Sahdev Zala (spzala@us.ibm.com), IBM

4327 Shitao li (lishitao@huawei.com), Huawei

4328 Simeon Monov (sdmonov@us.ibm.com), IBM

4329 Sivan Barzily, (sivan@gigaspaces.com), Gigaspaces

4330 Sridhar Ramaswamy (sramasw@brocade.com), Brocade

4331 Stephane Maes (stephane.maes@hp.com), HP

4332 Thomas Spatzier (thomas.spatzier@de.ibm.com), IBM

4333 Ton Ngo (ton@us.ibm.com), IBM

4334 Travis Tripp (travis.tripp@hp.com), HP

4335 Vahid Hashemian (vahidhashemian@us.ibm.com), IBM

4336 Wayne Witzel (wayne.witzel@canonical.com), Canonical

4337 Yaron Parasol (yaropa@gigaspaces.com), Gigaspaces

Appendix C. Revision History

Revision	Date	Editor	Changes Made
WD01, Rev01	2016-01-06	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • Initial WD01, Revision 01 baseline for TOSCA Simple Profile in YAML v1.1 • Cha. 10 Removed URL column for use cases in favor of a single link to Git directory where they can be found. • Metadata added to top-level entities • Policy grammar/schema fully defined. • Ch5. Defined TOSCA Entity Root type which is now the parent type for all TOSCA top-level types (i.e., Artifact, Capability, Relationship, Node, Group, Policy, etc.). Updated all top-level definitions to reflect in "derived_from" keyname. • Added TimeInterval Data Type • 3.5.16.1: Added keyname "schedule".
WD01, Rev02	2016-01-25	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • 5: Removed toasca.Root type from chapter 5 until ad-hoc can agree on use cases likely to come from the TOSCA instance model WG. • Cleaned up TOSCA Entity Root Reorganization.
WD01, Rev03	2016-03-22	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • 3.5.7, 3.9.3: Fixed "import" grammar (section 3.5.7) and reference to it in repository example (section 3.9.3.9.3) • 3.6.11.2 – Group Type – clarified group types could have members that were other groups types. • 5.2.5: Fixed NetworkInfo (section 5.2.5) example which was missing the 'properties' keyword. • 5.2.6: Clarified PortDef examples (section 5.2.6) • 5.2.7: Fixed PortSpec (section 5.2.7) definition to assure that target, target_range, source and source_range properties were not 'required' in schema. • • Fixed the following issues raised by TC Admin.: <ul style="list-style-type: none"> • <i>Margins should be 1" top, 1" right and left, 0.5" bottom. [this will center the "new" footer, which is currently offset].</i> • <i>The footer uses different font size (Arial 10 instead of Arial 8) and wording ("Standards Track Draft" instead of "Standards Track Work Product").</i> • <i>Set the following three styles to use Arial 10:</i> <ul style="list-style-type: none"> • "Normal around table" • "List Paragraph" • "List Bullet 3" • <i>Around section 2.10.1, we corrected some text in the wrong font by re-applying the "normal" style.</i> • <i>In Section 1.8 Glossary that "Node Template" definition starts off with "Relationship Template" Is that correct? Also, the paragraph formatting of the definitions seems to use weird indenting.</i> • <i>In section 5.7.4.4, the diagram overlays the footer. We fixed this on our side by setting the preceding paragraph attribute to "keep with next".</i> • <i>In section 2.10.2, second paragraph after Figure 1, there is a reference to "Section 0". The link jumps to 2.9.2. Is this correct?</i> • <i>The table of examples is labelled Table of Figures. Also, the paragraph styles of these two titles should be changed from "Body text" to "Level 1", so they will show up in the TOC.</i> • 3.6.5 Interface Type – missing "derived_from" in keyname table, grammar and example.

WDO1, Rev04	2016-03-23	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • 5.2: Added section discussing TOSCA normative type names, their treatment and requirements to respect case (i.e., be case sensitive) when processing. • 3.6: All data types that are entity types now have their keyname tables reference the common keynames listed in section 3.6.1 TOSCA Entity schema. • 3.6.11: Added attributes, requirements and capabilities keynames to Group Type making it more like a Node Type (no artifacts, still logical aggregator of a set of nodes). • 5.9.11: Added the “network” (i.e, Endpoint) and “storage” (i.e., Storage) capabilities to the Container.Application node type.
WDO1, Rev05	2016-04-20	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • 3.1: Bumped version number to 1.1 • 5.3.2: typo. ‘userh’ -> ‘user’ in keyname table • 3.6.4.4 Artifact Type - Added a note regarding “mime types” to reference official list of Apache mime types to give reader a suitable reference for expected values. •
WDO1, Rev06	2016-17-05	Luc Boutier, FastConnect	<ul style="list-style-type: none"> • 3.5.14.2.2 replaced Node Type by Node Template. • 3.5.17: Add workflow activity definition • 3.5.18: Add workflow precondition definition • 3.5.19: Add workflow step definition • 3.7.7.: Add Imperative workflow definition • 3.8: Add the workflows keyname to the topology template definition • 6.3: Added a simplified way to declare a CSAR without the meta file. • 7: Added a TOSCA workflows section.
WDO1, Rev07	2016-19-05	Luc Boutier, FastConnect	<ul style="list-style-type: none"> • 3.5.18: Add assertion definition • 3.5.19: Add condition clause definition • 3.5.20: Leverage condition clause in precondition definition • 3.5.21: Leverage condition clause as filter in step definition • 7.2: Add documentation and example on TOSCA normative weaving • 7.3: Fix examples in imperative workflows definition
WDO1, Rev08	2016-31-05	Luc Boutier, FastConnect	<ul style="list-style-type: none"> • 7.2: Specifies current expected declarative workflows and limitations.
WDO1, Rev09	2016-31-05	Luc Boutier, FastConnect; Matt Rutkowski, IBM	<ul style="list-style-type: none"> • 1.8: Add description for abstract nodes and no-op nodes to the glossary • Fixed typos, spelling/grammar and fixed numerous broken hyperlinks.
WDO1, post-CSD01	2016-07-11	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • 3.5.16 – invalid type for schema period. Correct in table (scalar-unit.time), incorrect in code schema listing (integer). • 3.1.3.1 – Added namespace collision requirements for policies and moved “groups” requirements to include types as well.

4339