



TOSCA Simple Profile in YAML Version 1.0

Committee Specification Draft 04 / Public Review Draft 01

27 August 2015

Specification URIs

This version:

<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/csprd01/TOSCA-Simple-Profile-YAML-v1.0-csprd01.pdf> (Authoritative)
<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/csprd01/TOSCA-Simple-Profile-YAML-v1.0-csprd01.html>
<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/csprd01/TOSCA-Simple-Profile-YAML-v1.0-csprd01.docx>

Previous version:

<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/csd03/TOSCA-Simple-Profile-YAML-v1.0-csd03.pdf> (Authoritative)
<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/csd03/TOSCA-Simple-Profile-YAML-v1.0-csd03.html>
<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/csd03/TOSCA-Simple-Profile-YAML-v1.0-csd03.doc>

Latest version:

<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/TOSCA-Simple-Profile-YAML-v1.0.pdf> (Authoritative)
<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/TOSCA-Simple-Profile-YAML-v1.0.html>
<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/TOSCA-Simple-Profile-YAML-v1.0.docx>

Technical Committee:

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

Chairs:

Paul Lipton (paul.lipton@ca.com), CA Technologies
Simon Moser (smoser@de.ibm.com), IBM

Editors:

Derek Palma (dpalma@vnomi.com), Vnomi
Matt Rutkowski (mrutkows@us.ibm.com), IBM
Thomas Spatzier (thomas.spatzier@de.ibm.com), IBM

Related work:

This specification is related to:

- *Topology and Orchestration Specification for Cloud Applications Version 1.0*. 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/1.0>

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

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>).

Citation format:

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

[TOSCA-Simple-Profile-YAML-v1.0]

TOSCA Simple Profile in YAML Version 1.0. Edited by Derek Palma, Matt Rutkowski, and Thomas Spatzier. 27 August 2015. OASIS Committee Specification Draft 04 / Public Review Draft 01. <http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/csprd01/TOSCA-Simple-Profile-YAML-v1.0-csprd01.html>. Latest version: <http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/TOSCA-Simple-Profile-YAML-v1.0.html>.

Notices

Copyright © OASIS Open 2015. 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	10
Table of Figures	10
1 Introduction	11
1.1 Objective	11
1.2 Summary of key TOSCA concepts	11
1.3 Implementations	11
1.4 Terminology	12
1.5 Notational Conventions	12
1.6 Normative References	12
1.7 Non-Normative References	12
1.8 Glossary	12
2 TOSCA by example	14
2.1 A “hello world” template for TOSCA Simple Profile in YAML	14
2.1.1 Requesting input parameters and providing output	15
2.2 TOSCA template for a simple software installation	16
2.3 Overriding behavior of predefined node types	18
2.4 TOSCA template for database content deployment	19
2.5 TOSCA template for a two-tier application	21
2.6 Using a custom script to establish a relationship in a template	24
2.7 Using custom relationship types in a TOSCA template	25
2.7.1 Definition of a custom relationship type	26
2.8 Defining generic dependencies between nodes in a template	26
2.9 Describing abstract requirements for nodes and capabilities in a TOSCA template	27
2.9.1 Using a node_filter to define hosting infrastructure requirements for a software	28
2.9.2 Using an abstract node template to define infrastructure requirements for software	29
2.9.3 Using a node_filter to define requirements on a database for an application	30
2.10 Using node template substitution for model composition	31
2.10.1 Understanding node template instantiation through a TOSCA Orchestrator	32
2.10.2 Definition of the top-level service template	32
2.10.3 Definition of the database stack in a service template	33
2.11 Using node template substitution for chaining subsystems	35
2.11.1 Defining the overall subsystem chain	36
2.11.2 Defining a subsystem (node) type	37
2.11.3 Defining the details of a subsystem	38
2.12 Grouping node templates	41
2.13 Using YAML Macros to simplify templates	43
2.14 Passing information as inputs to Nodes and Relationships	44
2.14.1 Example: declaring input variables for all operations on a single interface	44
2.14.2 Example: declaring input variables for a single operation	45
2.14.3 Example: setting output variables to an attribute	45
2.14.4 Example: passing output variables between operations	45
2.15 Topology Template Model versus Instance Model	46
2.16 Using attributes implicitly reflected from properties	46

3	TOSCA Simple Profile definitions in YAML	48
3.1	TOSCA Namespace URI and alias	48
3.1.1	TOSCA Namespace prefix	48
3.1.2	TOSCA Namespacing in TOSCA Service Templates	48
3.1.3	Rules to avoid namespace collisions	48
3.2	Parameter and property types	49
3.2.1	Referenced YAML Types	49
3.2.2	TOSCA version.....	50
3.2.3	TOCSA range type	51
3.2.4	TOSCA list type	52
3.2.5	TOSCA map type	53
3.2.6	TOCSA scalar-unit type.....	55
3.3	Normative values	58
3.3.1	Node States.....	58
3.3.2	Relationship States	58
3.3.3	Directives.....	59
3.3.4	Network Name aliases	59
3.4	TOSCA Metamodel.....	59
3.4.1	Required Keynames.....	59
3.5	Reusable modeling definitions.....	59
3.5.1	Description definition	59
3.5.2	Constraint clause.....	60
3.5.3	Property Filter definition	62
3.5.4	Node Filter definition	63
3.5.5	Repository definition.....	65
3.5.6	Artifact definition	66
3.5.7	Import definition	67
3.5.8	Property definition	68
3.5.9	Property assignment	70
3.5.10	Attribute definition.....	71
3.5.11	Attribute assignment.....	72
3.5.12	Parameter definition	73
3.5.13	Operation definition	75
3.5.14	Interface definition	77
3.6	Type-specific definitions	78
3.6.1	Capability definition	78
3.6.2	Requirement definition	80
3.6.3	Artifact Type	82
3.6.4	Interface Type.....	83
3.6.5	Data Type	85
3.6.6	Capability Type.....	86
3.6.7	Requirement Type.....	88
3.6.8	Node Type	88
3.6.9	Relationship Type.....	90
3.6.10	Group Type.....	91

3.6.11 Policy Type	93
3.7 Template-specific definitions	94
3.7.1 Capability assignment	94
3.7.2 Requirement assignment	95
3.7.3 Node Template	99
3.7.4 Relationship Template	101
3.7.5 Group definition	102
3.7.6 Policy definition	103
3.8 Topology Template definition	104
3.8.1 Keynames	104
3.8.2 Grammar	105
3.9 Service Template definition	110
3.9.1 Keynames	110
3.9.2 Grammar	111
3.9.3 Top-level keyname definitions	112
4 TOSCA functions	122
4.1 Reserved Function Keywords	122
4.2 Environment Variable Conventions	122
4.2.1 Reserved Environment Variable Names and Usage	122
4.2.2 Prefixed vs. Unprefixed TARGET names	124
4.3 Intrinsic functions	124
4.3.1 concat	125
4.3.2 token	125
4.4 Property functions	126
4.4.1 get_input	126
4.4.2 get_property	126
4.5 Attribute functions	128
4.5.1 get_attribute	128
4.6 Operation functions	129
4.6.1 get_operation_output	129
4.7 Navigation functions	130
4.7.1 get_nodes_of_type	130
4.8 Artifact functions	130
4.8.1 get_artifact	130
4.9 Context-based Entity names (global)	133
5 TOSCA normative type definitions	134
5.1 Assumptions	134
5.2 Data Types	134
5.2.1 tosa.datatypes.Root	134
5.2.2 tosa.datatypes.Credential	134
5.2.3 tosa.datatypes.network.NetworkInfo	136
5.2.4 tosa.datatypes.network.PortInfo	137
5.2.5 tosa.datatypes.network.PortDef	139
5.2.6 tosa.datatypes.network.PortSpec	139
5.3 Artifact Types	141

5.3.1	tosca.artifacts.Root.....	141
5.3.2	tosca.artifacts.File	141
5.3.3	Deployment Types	141
5.3.4	Implementation Types	142
5.4	Capabilities Types	144
5.4.1	tosca.capabilities.Root	144
5.4.2	tosca.capabilities.Node	144
5.4.3	tosca.capabilities.Container	144
5.4.4	tosca.capabilities.Endpoint	145
5.4.5	tosca.capabilities.Endpoint.Public	147
5.4.6	tosca.capabilities.Endpoint.Admin	148
5.4.7	tosca.capabilities.Endpoint.Database	148
5.4.8	tosca.capabilities.Attachment	149
5.4.9	tosca.capabilities.OperatingSystem	149
5.4.10	tosca.capabilities.Scalable	150
5.4.11	tosca.capabilities.network.Bindable	151
5.5	Requirement Types	152
5.6	Relationship Types	152
5.6.1	tosca.relationships.Root	152
5.6.2	tosca.relationships.DependsOn	153
5.6.3	tosca.relationships.HostedOn	153
5.6.4	tosca.relationships.ConnectsTo	153
5.6.5	tosca.relationships.AttachesTo	154
5.6.6	tosca.relationships.RoutesTo	155
5.7	Interface Types	155
5.7.1	Additional Requirements	155
5.7.2	Best Practices.....	155
5.7.3	tosca.interfaces.node.lifecycle.Standard.....	155
5.7.4	tosca.interfaces.relationship.Configure	158
5.8	Node Types.....	160
5.8.1	tosca.nodes.Root	160
5.8.2	tosca.nodes.Compute	161
5.8.3	tosca.nodes.SoftwareComponent	163
5.8.4	tosca.nodes.WebServer	164
5.8.5	tosca.nodes.WebApplication	165
5.8.6	tosca.nodes.DBMS.....	165
5.8.7	tosca.nodes.Database.....	166
5.8.8	tosca.nodes.ObjectStorage	167
5.8.9	tosca.nodes.BlockStorage.....	168
5.8.10	tosca.nodes.Container.Runtime	170
5.8.11	tosca.nodes.Container.Application.....	170
5.8.12	tosca.nodes.LoadBalancer	170
5.9	Group Types	171
5.9.1	tosca.groups.Root	171
5.10	Policy Types	172

5.10.1	tosca.policies.Root	172
5.10.2	tosca.policies.Placement.....	172
5.10.3	tosca.policies.Scaling	172
5.10.4	tosca.policies.Update	173
5.10.5	tosca.policies.Performance	173
6	TOSCA Cloud Service Archive (CSAR) format.....	174
6.1	Overall Structure of a CSAR.....	174
6.2	TOSCA Meta File.....	174
6.2.1	Example.....	174
7	TOSCA networking.....	175
7.1	Networking and Service Template Portability.....	175
7.2	Connectivity Semantics	175
7.3	Expressing connectivity semantics.....	176
7.3.1	Connection initiation semantics.....	176
7.3.2	Specifying layer 4 ports.....	177
7.4	Network provisioning	178
7.4.1	Declarative network provisioning.....	178
7.4.2	Implicit network fulfillment	178
7.4.3	Controlling network fulfillment	178
7.5	Network Types	181
7.5.1	tosca.nodes.network.Network	181
7.5.2	tosca.nodes.network.Port.....	183
7.5.3	tosca.capabilities.network.Linkable.....	185
7.5.4	tosca.relationships.network.LinksTo	185
7.5.5	tosca.relationships.network.BindsTo.....	186
7.6	Network modeling approaches	186
7.6.1	Option 1: Specifying a network outside the application's Service Template.....	186
7.6.2	Option 2: Specifying network requirements within the application's Service Template	190
8	Non-normative type definitions.....	192
8.1	Artifact Types	192
8.1.1	tosca.artifacts.Deployment.Image.Container.Docker.....	192
8.1.2	tosca.artifacts.Deployment.Image.VM.ISO	192
8.1.3	tosca.artifacts.Deployment.Image.VM.QCOW2.....	192
8.2	Capability Types	192
8.2.1	tosca.capabilities.Container.Docker	192
8.3	Node Types.....	194
8.3.1	tosca.nodes.Database.MySQL.....	194
8.3.2	tosca.nodes.DBMS.MySQL.....	194
8.3.3	tosca.nodes.WebServer.Apache	195
8.3.4	tosca.nodes.WebApplication.WordPress	195
8.3.5	tosca.nodes.WebServer.Nodejs.....	196
8.3.6	tosca.nodes.Container.Application.Docker	196
9	Component Modeling Use Cases.....	197
9.1.1	Use Case: Exploring the HostedOn relationship using WebApplication and WebServer.....	197
9.1.2	Use Case: Establishing a ConnectsTo relationship to WebServer	198

9.1.3 Use Case: Attaching (local) BlockStorage to a Compute node	199
9.1.4 Use Case: Reusing a BlockStorage Relationship using Relationship Type or Relationship Template.....	199
10 Application Modeling Use Cases.....	204
10.1 Use cases	204
10.1.1 Overview.....	204
10.1.2 Compute: Create a single Compute instance with a host Operating System	206
10.1.3 Software Component 1: Automatic deployment of a Virtual Machine (VM) image artifact ...	208
10.1.4 Block Storage 1: Using the normative AttachesTo Relationship Type	210
10.1.5 Block Storage 2: Using a custom AttachesTo Relationship Type.....	212
10.1.6 Block Storage 3: Using a Relationship Template of type AttachesTo	215
10.1.7 Block Storage 4: Single Block Storage shared by 2-Tier Application with custom AttachesTo Type and implied relationships.....	217
10.1.8 Block Storage 5: Single Block Storage shared by 2-Tier Application with custom AttachesTo Type and explicit Relationship Templates.....	219
10.1.9 Block Storage 6: Multiple Block Storage attached to different Servers	222
10.1.10 Object Storage 1: Creating an Object Storage service	225
10.1.11 Network 1: Server bound to a new network	226
10.1.12 Network 2: Server bound to an existing network.....	228
10.1.13 Network 3: Two servers bound to a single network	229
10.1.14 Network 4: Server bound to three networks.....	231
10.1.15 WebServer-DBMS 1: WordPress + MySQL, single instance.....	234
10.1.16 WebServer-DBMS 2: Nodejs with PayPal Sample App and MongoDB on separate instances	238
10.1.17 Multi-Tier-1: Elasticsearch, Logstash, Kibana (ELK) use case with multiple instances	242
10.1.18 Container-1: Containers using Docker single Compute instance (Containers only)	248
11 TOSCA Policies.....	250
11.1 A declarative approach	250
11.1.1 Declarative considerations	250
11.2 Consideration of Event, Condition and Action	250
11.3 Types of policies	250
11.3.1 Access control policies	250
11.3.2 Placement policies	251
11.3.3 Quality-of-Service (QoS) policies	251
11.4 Policy relationship considerations	251
11.5 Use Cases	252
11.5.1 Placement.....	252
11.5.2 Scaling.....	253
12 Conformance	255
12.1 Conformance Targets	255
12.2 Conformance Clause 1: TOSCA YAML service template	255
12.3 Conformance Clause 2: TOSCA processor.....	255
12.4 Conformance Clause 3: TOSCA orchestrator	255
12.5 Conformance Clause 4: TOSCA generator	256
12.6 Conformance Clause 5: TOSCA archive	256
Appendix A. Known Extensions to TOSCA v1.0.....	257

A.1 Model Changes	257
A.2 Normative Types	257
Appendix B. Acknowledgments	259
Appendix C. Revision History.....	260

Table of Examples

Example 1 - TOSCA Simple "Hello World"	14
Example 2 - Template with input and output parameter sections	15
Example 3 - Simple (MySQL) software installation on a TOSCA Compute node	16
Example 4 - Node Template overriding its Node Type's "configure" interface	18
Example 5 - Template for deploying database content on-top of MySQL DBMS middleware	19
Example 6 - Basic two-tier application (web application and database server tiers)	21
Example 7 - Providing a custom relationship script to establish a connection	24
Example 8 - A web application Node Template requiring a custom database connection type	25
Example 9 - Defining a custom relationship type	26
Example 10 - Simple dependency relationship between two nodes.....	27
Example 11 - An abstract "host" requirement using a node filter	28
Example 12 - An abstract Compute node template with a node filter.....	29
Example 13 - An abstract database requirement using a node filter.....	30
Example 14 - An abstract database node template	31
Example 15 - Referencing an abstract database node template.....	33
Example 16 - Using substitution mappings to export a database implementation	34
Example 17 - Declaring a transaction subsystem as a chain of substitutable node templates	36
Example 18 - Defining a TransactionSubsystem node type.....	38
Example 19 - Implementation of a TransactionSubsystem node type using substitution mappings	39
Example 20 - Grouping Node Templates for possible policy application	41
Example 21 - Grouping nodes for anti-colocation policy application.....	42

Table of Figures

Figure 1: Using template substitution to implement a database tier.....	32
Figure 2: Substitution mappings	34
Figure 3: Chaining of subsystems in a service template.....	36
Figure 4: Defining subsystem details in a service template.....	39
Figure-5: Typical 3-Tier Network.....	178
Figure-6: Generic Service Template.....	187
Figure-7: Service template with network template A	187
Figure-8: Service template with network template B	188

1 Introduction

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

1.3 Implementations

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

- TOSCA YAML service template: A YAML document artifact containing a (TOSCA) service template that represents a Cloud application.
- TOSCA processor: An engine or tool that is capable of parsing and interpreting a TOSCA YAML service template for a particular purpose. For example, the purpose could be validation, translation or visual rendering.
- TOSCA orchestrator (also called orchestration engine): A TOSCA processor that interprets a TOSCA YAML service template then instantiates and deploys the described application in a Cloud.

- TOSCA generator: A tool that generates a TOSCA YAML service template. An example of generator is a modeling tool capable of generating or editing a TOSCA YAML service template (often such a tool would also be a TOSCA processor).
- TOSCA archive (or TOSCA Cloud Service Archive, or “CSAR”): a package artifact that contains a TOSCA YAML service template and other artifacts usable by a TOSCA orchestrator to deploy an application.

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

1.4 Terminology

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

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

1.5 Notational Conventions

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

1.6 Normative References

[RFC2119]	S. Bradner, <i>Key words for use in RFCs to Indicate Requirement Levels</i> , http://www.ietf.org/rfc/rfc2119.txt , IETF RFC 2119, March 1997.
[TOSCA-1.0]	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

1.7 Non-Normative References

[AWS-CFN]	Amazon Cloud Formation (CFN), http://aws.amazon.com/cloudformation/
[Chef]	Chef, https://wiki.opscode.com/display/chef/Home
[OS-Heat]	OpenStack Project Heat, https://wiki.openstack.org/wiki/Heat
[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

1.8 Glossary

The following terms are used throughout this specification and have the following definitions when used in 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>Relationship 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	<p>A Service Template 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.</p> <p>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.</p>
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	<p>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.</p> <p>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.</p>

2 TOSCA by example

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

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

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

Example 1 - TOSCA Simple "Hello World"

```
tosca_definitions_version: tosca_simple_yaml_1_0

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

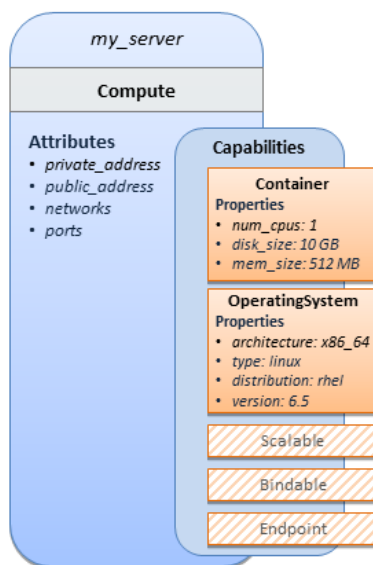
topology_template:
  node_templates:
    my_server:
      type: tosca.nodes.Compute
      capabilities:
        # Host container properties
        host:
          properties:
            num_cpus: 1
            disk_size: 10 GB
            mem_size: 4096 MB
        # Guest Operating System properties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

The template above contains a very simple topology template with only a single ‘Compute’ node template that declares some basic values for properties within two of the several capabilities that are built into the Compute node type definition. All TOSCA Orchestrators are expected to know how to instantiate a Compute node since it is normative and expected to represent a well-known function that is portable across TOSCA implementations. This expectation is true for all normative TOSCA Node and Relationship types that are defined in the Simple Profile specification. This means, with TOSCA’s approach, that the application developer does not need to provide any deployment or implementation artifacts that contain code or logic to orchestrate these common

89 software components. TOSCA orchestrators simply select or allocate the correct node (resource) type that fulfils
 90 the application topologies requirements using the properties declared in the node and its capabilities.

91 In the above example, the “**host**” capability contains properties that allow application developers to optionally
 92 supply the number of CPUs, memory size and disk size they believe they need when the Compute node is
 93 instantiated in order to run their applications. Similarly, the “**os**” capability is used to provide values to indicate
 94 what host operating system the Compute node should have when it is instantiated.

95 The logical diagram of the “hello world” Compute node would look as follows:



96
 97
 98 As you can see, the **Compute** node also has attributes and other built-in capabilities, such as **Bindable** and
 99 **Endpoint**, each with additional properties that will be discussed in other examples later in this document.
 100 Although the Compute node has no direct properties apart from those in its capabilities, other TOSCA node type
 101 definitions may have properties that are part of the node type itself in addition to having Capabilities. TOSCA
 102 orchestration engines are expected to validate all property values provided in a node template against the
 103 property definitions in their respective node type definitions referenced in the service template. The
 104 **tosca_definitions_version** keyname in the TOSCA service template identifies the versioned set of normative
 105 TOSCA type definitions to use for validating those types defined in the TOSCA Simple Profile including the
 106 Compute node type. Specifically, the value **tosca_simple_yaml_1_0** indicates Simple Profile v1.0.0 definitions
 107 would be used for validation. Other type definitions may be imported from other service templates using the
 108 **import** keyword discussed later.

109 2.1.1 Requesting input parameters and providing output

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

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

```
tosca_definitions_version: tosa_simple_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 ] }

```

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

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

123 2.2 TOSCA template for a simple software installation

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

128 *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

```

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

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

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



141

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

151 2.3 Overriding behavior of predefined node types

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

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

161 *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
```

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

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

169 2.4 TOSCA template for database content deployment

170 In the example shown in section 4 the deployment of the MySQL middleware only, i.e. without actual database
171 content was shown. The following example shows how such a template can be extended to also contain the
172 definition of custom database content on-top of the MySQL DBMS software.

173 *Example 5 - Template for deploying database content on-top of MySQL DBMS middleware*

```
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: tosa.nodes.DBMS.MySQL
properties:
  root_password: { get_input: mysql_rootpw }
  port: { get_input: mysql_port }
requirements:
  - host: db_server

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

```

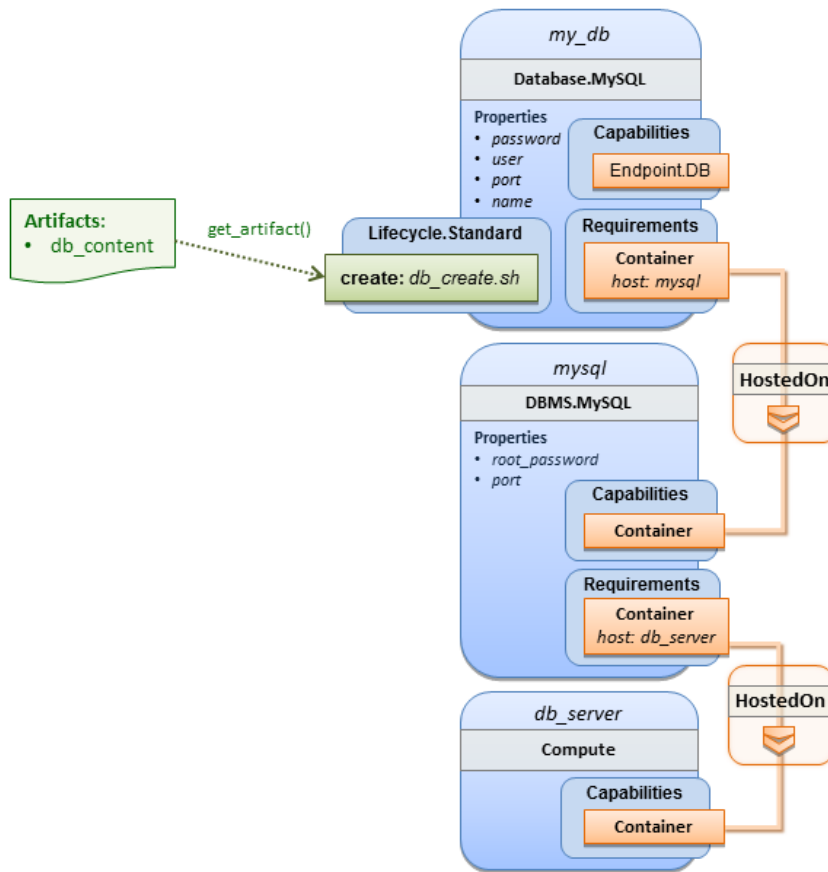
174 In the example above, the **my_db** node template or type **tosa.nodes.Database.MySQL** represents an actual
 175 MySQL database instance managed by a MySQL DBMS installation. The **requirements** section of the **my_db**
 176 node template expresses that the database it represents is to be hosted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

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

182 As you can see above, a script is associated with the create operation with the name **db_create.sh**. The
 183 TOSCA Orchestrator sees that this is not a named artifact declared in the node's artifact section, but instead a
 184 filename for a normative TOCA implementation artifact script type (i.e.,
 185 **tosa.artifacts.Implementation.Bash**). Since this is an implementation type for TOSCA, the orchestrator
 186 will execute the script automatically to create the node on **db_server**, but first it will prepare the local environment
 187 with the declared inputs for the operation. In this case, the orchestrator would see that the **db_data** input is using
 188 the **get_artifact** function to retrieve the file (**my_db_content.txt**) which is associated with the **db_content**
 189 artifact name prior to executing the **db_create.sh** script.

190

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



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

2.5 TOSCA template for a two-tier application

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

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: tosa.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: tosa.nodes.WebServer.Apache
    properties:
      # omitted here for brevity
    requirements:
      - host: web_server

```

```

web_server:
  type: tosa.nodes.Compute
  capabilities:
    # omitted here for brevity

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

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

```

206 The web application stack consists of the **wordpress**, the **apache** and the **web_server** node templates. The
 207 **wordpress** node template represents a custom web application of type
 208 **tosa.nodes.WebApplication.WordPress** which is hosted on an Apache web server represented by the **apache**
 209 node template. This hosting relationship is expressed via the **host** entry in the **requirements** section of the
 210 **wordpress** node template. The **apache** node template, finally, is hosted on the **web_server** compute node.

211 The database stack consists of the **wordpress_db**, the **mysql** and the **db_server** node templates. The
 212 **wordpress_db** node represents a custom database of type **tosa.nodes.Database.MySQL** which is hosted on a
 213 MySQL DBMS represented by the **mysql** node template. This node, in turn, is hosted on the **db_server** compute
 214 node.

215 The **wordpress** node requires a connection to the **wordpress_db** node, since the WordPress application needs
 216 a database to store its data in. This relationship is established through the **database_endpoint** entry in the
 217 **requirements** section of the **wordpress** node template's declared node type. For configuring the WordPress web
 218 application, information about the database to connect to is required as input to the **configure** operation.
 219 Therefore, the input parameters are defined and values for them are retrieved from the properties and attributes
 220 of the **wordpress_db** node via the **get_property** and **get_attribute** functions. In the above example, these
 221 inputs are defined at the interface-level and would be available to all operations of the **Standard** interface (i.e.,

222 the `tosca.interfaces.node.lifecycle.Standard` interface) within the `wordpress` node template and not just
223 the `configure` operation.

224 2.6 Using a custom script to establish a relationship in a template

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

230 In some cases it might be necessary to provide special processing logic to be executed when establishing
231 relationships between nodes at runtime. For example, when connecting the WordPress application from previous
232 examples to the MySQL database, it might be desired to apply custom configuration logic in addition to that
233 already implemented in the application node type. In such a case, it is possible for the template author to provide
234 a custom script as implementation for an operation to be executed at runtime as shown in the following example.

235 *Example 7 – Providing a custom relationship script to establish a connection*

```
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_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 ...

```

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

The ConnectsTo relationship (see 5.6.4) also provides a default **Configure** interface with operations that optionally get executed when the orchestrator establishes the relationship. In the above example, the author has provided the custom script **wp_db_configure.sh** to be executed for the operation called **pre_configure_source**. The script file is assumed to be located relative to the referencing service template such as a relative directory within the TOSCA Cloud Service Archive (CSAR) packaging format. This approach allows for conveniently hooking in custom behavior without having to define a completely new derived relationship type.

2.7 Using custom relationship types in a TOSCA template

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

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

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 ...

```

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

2.7.1 Definition of a custom relationship type

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

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

```

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

2.8 Defining generic dependencies between nodes in a template

In some cases it can be necessary to define a generic dependency between two nodes in a template to influence orchestration behavior, i.e. to first have one node processed before another dependent node gets processed. This can be done by using the generic **dependency** requirement which is defined by the [TOSCA Root Node Type](#) and thus gets inherited by all other node types in TOSCA (see section 5.8.1).

```

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

```

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

284 2.9 Describing abstract requirements for nodes and capabilities in a 285 TOSCA template

286 In TOSCA templates, nodes are either:

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

291
 292 TOSCA Orchestrators, by default, when finding an abstract node in TOSCA Service Template during deployment
 293 will attempt to “select” a concrete implementation for the abstract node type that best matches and fulfills the
 294 requirements and property constraints the template author provided for that abstract node. The concrete
 295 implementation of the node could be provided by another TOSCA Service Template (perhaps located in a catalog
 296 or repository known to the TOSCA Orchestrator) or by an existing resource or service available within the target
 297 Cloud Provider’s platform that the TOSCA Orchestrator already has knowledge of.

298
 299 TOSCA supports two methods for template authors to express requirements for an abstract node within a TOSCA
 300 service template.

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

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

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

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

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

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

```

324 In the example above, the **mysql** component contains a **host** requirement for a node of type **Compute** which it
 325 inherits from its parent DBMS node type definition; however, there is no declaration or reference to any node
 326 template of type **Compute**. Instead, the **mysql** node template augments the abstract "**host**" requirement with a
 327 **node_filter** which contains additional selection criteria (in the form of property constraints that the provider
 328 must use when selecting or allocating a host **Compute** node.

329 Some of the constraints shown above narrow down the boundaries of allowed values for certain properties such
 330 as **mem_size** or **num_cpus** for the "**host**" capability by means of qualifier functions such as **greater_or_equal**.
 331 Other constraints, express specific values such as for the **architecture** or **distribution** properties of the "**os**"
 332 capability which will require the provider to find a precise match.

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

335 2.9.2 Using an abstract node template to define infrastructure requirements for 336 software

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

342 *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: tosca.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

```

343 As you can see the resulting `mysql_compute` node template looks very much like the “hello world” template as
 344 shown in Chapter 2 (where the `Compute` node template was abstract), but this one also allows the TOSCA
 345 orchestrator more flexibility when “selecting” a host `Compute` node by providing flexible constraints for properties
 346 like `mem_size`.

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

352 2.9.3 Using a `node_filter` to define requirements on a database for an application

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

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

```

tosca_definitions_version: tosca_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 }

```

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

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

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

369 *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 }
```

370 2.10 Using node template substitution for model composition

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

2.10.1 Understanding node template instantiation through a TOSCA Orchestrator

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

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

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

2.10.2 Definition of the top-level service template

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

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

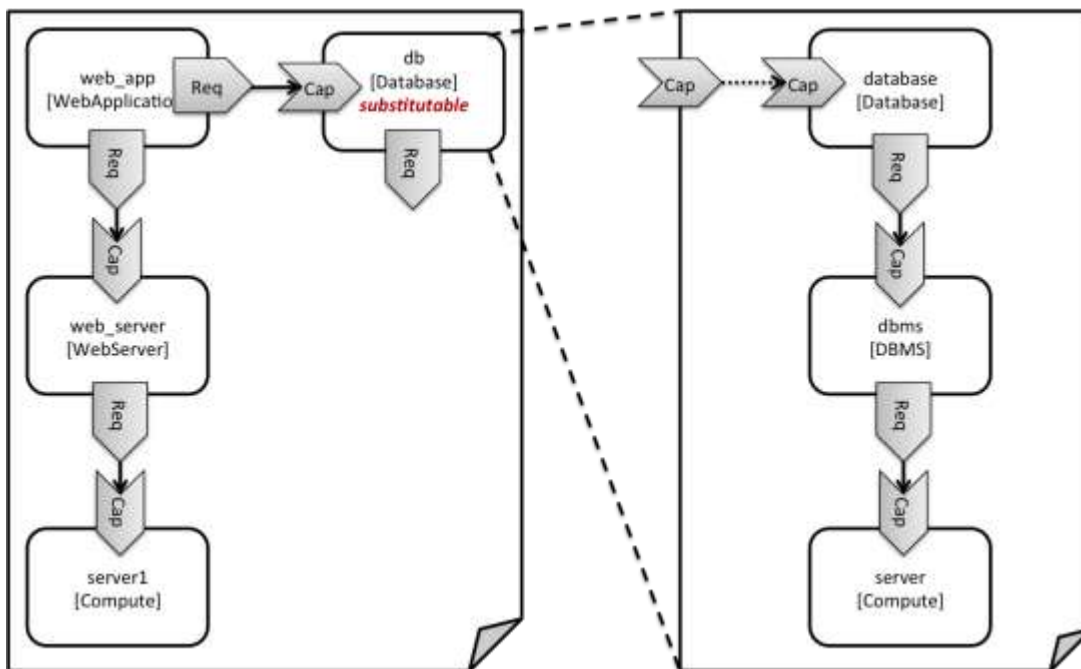


Figure 1: Using template substitution to implement a database tier

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

408 Note that in contrast to the use case described in section 0 (where a database was abstractly referred to in the
409 **requirements** section of a node and the database itself was not represented as a node template), the approach
410 shown here allows for some additional modeling capabilities in cases where this is required.
411 For example, if multiple components shall use the same database (or any other sub-system of the overall
412 service), this can be expressed by means of normal relations between node templates, whereas such modeling
413 would not be possible in **requirements** sections of disjoint node templates.

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

```
tosca_definitions_version: tosca_simple_yaml_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: tosca.nodes.WebServer
      requirements:
        - host: server

    server:
      type: tosca.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: tosca.nodes.Database
      properties:
        user: my_db_user
        password: secret
        name: my_db_name
```

415 2.10.3 Definition of the database stack in a service template

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

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

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

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

- 428 1. It defines what node templates, i.e. node templates of which type, can be substituted by the topology
 429 template.
- 430 2. It defines how capabilities of the substituted node (or the capabilities defined by the node type of the
 431 substituted node template, respectively) are bound to capabilities of node templates defined in the
 432 topology template.
- 433 3. It defines how requirements of the substituted node (or the requirements defined by the node type of the
 434 substituted node template, respectively) are bound to requirements of node templates defined in the
 435 topology template.

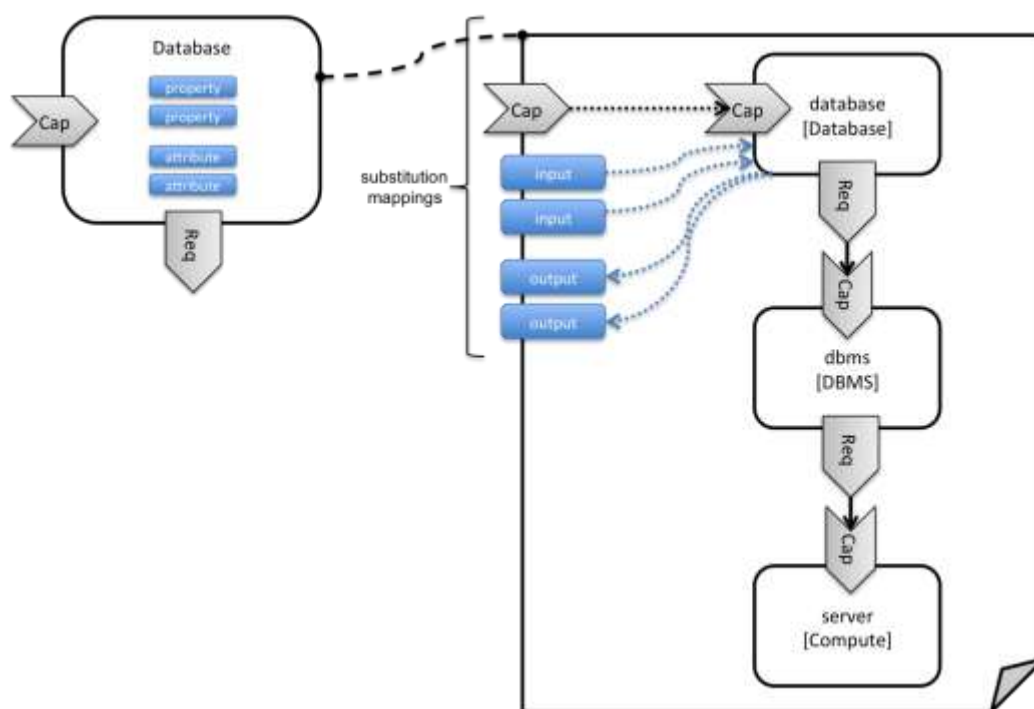


Figure 2: Substitution mappings

436
 437 The **substitution_mappings** section in the sample below denotes that this topology template can be used for
 438 substituting node templates of type **tosca.nodes.Database**. It further denotes that the **database_endpoint**
 439 capability of the substituted node gets fulfilled by the **database_endpoint** capability of the **database** node
 440 contained in the topology template.
 441

442 *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

```

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

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

451 2.11 Using node template substitution for chaining subsystems

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

2.11.1 Defining the overall subsystem chain

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

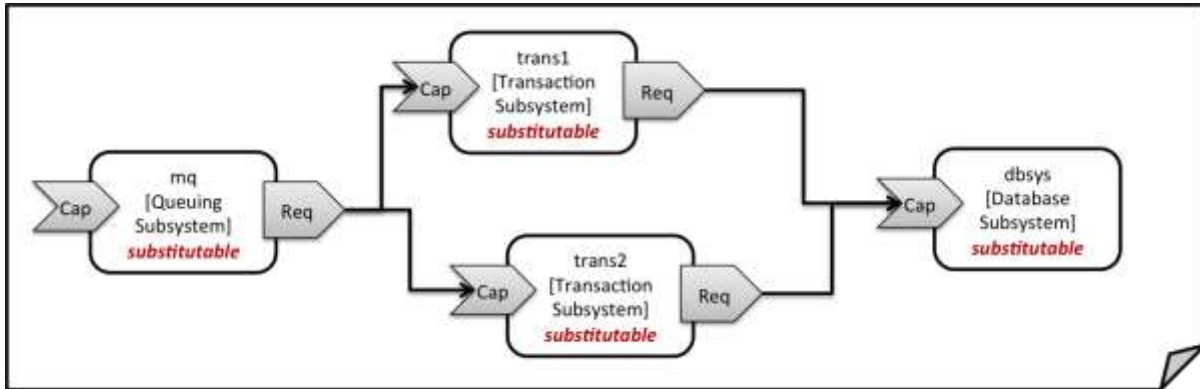


Figure 3: Chaining of subsystems in a service template

All of the node templates in the service template shown above are abstract and considered substitutable where each can be treated as their own subsystem; therefore, when instantiating the overall service, the orchestrator would realize each substitutable node template using other TOSCA service templates. These service templates would include more nodes and relationships that include the details for each subsystem. A simplified version of a TOSCA service template for the overall service is given in the following listing.

Example 17 – Declaring a transaction subsystem as a chain of substitutable node templates

```
tosca_definitions_version: tosca_simple_yaml_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

```

479

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

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

487 If attributes of the instantiated subsystems shall be obtained, this would be possible by using the **get_attribute**
 488 intrinsic function on the respective subsystem node templates.

489 2.11.2 Defining a subsystem (node) type

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

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

495

```
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
```

497

498 Configuration parameters that shall be allowed for customizing the instantiation of any subsystem are defined as
499 properties of the node type. In the current example, those are the properties **mq_service_ip** and **receiver_port**
500 that had been used in the end-to-end service template in section 2.11.1.

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

504 2.11.3 Defining the details of a subsystem

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

510

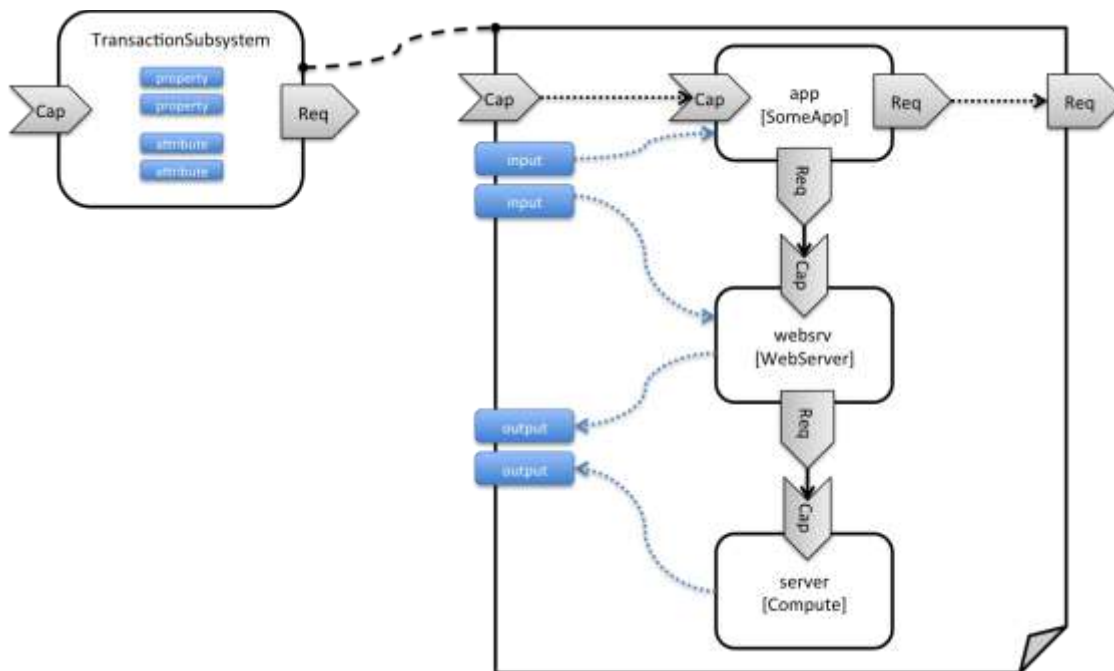


Figure 4: Defining subsystem details in a service template

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

The application named **app** provides a capability to receive messages, which is bound to the **message_receiver** capability of the substitutable node type. It further requires access to a database, so the application's **database_endpoint** requirement is mapped to the **database_endpoint** requirement of the **TransactionSubsystem** node type.

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

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

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

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 ] }

```

2.12 Grouping node templates

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

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

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: toska.nodes.WebServer.Apache
      properties:
        # Details omitted for brevity
      requirements:
        - host: server

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

  groups:
    webserver_group:
      type: toska.groups.Root

```

```
members: [ apache, server ]
```

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

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

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

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

557 *Example 21 - Grouping nodes for anti-colocation 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: tosa.nodes.DBMS.MySQL
properties:
  # omitted here for brevity
requirements:
  - host:
      node: tosa.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: tosa.groups.Root
    targets: [ 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

```

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

562 2.13 Using YAML Macros to simplify templates

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

568
 569 For example, an application that has a web server and database (i.e., a two-tier application) may be described
 570 using two **Compute** nodes (one to host the web server and another to host the database). The author may want
 571 both Compute nodes to be instantiated with similar properties such as operating system, distribution, version, etc.

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

```

tosca_definitions_version: tosca_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

```

2.14 Passing information as inputs to Nodes and Relationships

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

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

```

node_templates:
  wordpress:
    type: tosca.nodes.WebApplication.WordPress
    requirements:
      ...
      - database_endpoint: mysql_database
    interfaces:
      Standard:
        inputs:

```

```
wp_db_port: { get_property: [ SELF, database_endpoint, port ] }
```

2.14.2 Example: declaring input variables for a single operation

```
node_templates:
  wordpress:
    type: tosa.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 ] }
```

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

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
```

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

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 ] }
```

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

593 2.15 Topology Template Model versus Instance Model

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

604 2.16 Using attributes implicitly reflected from properties

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

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

612
613 Use of this feature is shown in the example below where a source node named **my_client**, of type **ClientNode**,
614 requires a connection to another node named **my_server** of type **ServerNode**. As you can see, the **ServerNode**
615 type defines a property named **notification_port** which defines a dedicated port number which instances of
616 **my_client** may use to post asynchronous notifications to it during runtime. In this case, the TOSCA Simple
617 Profile assures that the **notification_port** property is implicitly reflected as an attribute in the **my_server** node
618 (also with the name **notification_port**) when its node template is instantiated.
619

```
tosca_definitions_version: tosa_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

```

620

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

626

627 It should be noted that the actual port value of the **notification_port** attribute may or may not be the value
 628 **8000** as requested on the property; therefore, any node that is dependent on knowing its actual “runtime” value
 629 would use the **get_attribute** function instead of the **get_property** function.

3 TOSCA Simple Profile definitions in YAML

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

3.1 TOSCA Namespace URI and alias

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

Namespace Alias	Namespace URI	Specification Description
tosca_simple_yaml_1_0	http://docs.oasis-open.org/tosca/ns/simple/yaml/1.0	The TOSCA Simple Profile v1.0 (YAML) target namespace and namespace alias.

3.1.1 TOSCA Namespace prefix

The following TOSCA Namespace prefix is a reserved value and SHALL be used to reference the default 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

3.1.2 TOSCA Namespacing in TOSCA Service Templates

In the TOSCA Simple Profile, TOSCA Service Templates must always have, as the first line of YAML, the keyword “**tosca_definitions_version**” with an associated TOSCA Namespace Alias value. This single line accomplishes the following:

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

3.1.3 Rules to avoid namespace collisions

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

3.1.3.1 Additional Requirements

- Since TOSCA Service Templates can import (or substitute in) other Service Templates, TOSCA Orchestrators and tooling will encounter the “**tosca_definitions_version**” statement for each imported template. In these cases, the following additional requirements apply:

- Imported type definitions with the same Namespace URI, local name and version SHALL be equivalent.
 - If different values of the “**tosca_definitions_version**” are encountered, their corresponding type definitions MUST be uniquely identifiable using their corresponding Namespace URI using a different Namespace prefix.
- Duplicate local names (i.e., within the same Service Template SHALL be considered an error. These include, but are not limited to duplicate names found for the following definitions:
 - Repositories (repositories)
 - Data Types (data_types)
 - Node Types (node_types)
 - Relationship Types (relationship_types_
 - Capability Types (capability_types)
 - Artifact Types (artifact_types_
 - Interface Types (interface_types)
- Duplicate Template names within a Service Template’s Topology Template SHALL be considered an error. These include, but are not limited to duplicate names found for the following template types:
 - Node Templates (node_templates)
 - Relationship Templates (relationship_templates)
 - Inputs (inputs)
 - Outputs (outputs)
 - Groups (groups)
- Duplicate names for the following keynames within Types or Templates SHALL be considered an error. These include, but are not limited to duplicate names found for the following keynames:
 - Properties (properties)
 - Attributes (attributes)
 - Artifacts (artifacts)
 - Requirements (requirements)
 - Capabilities (capabilities)
 - Interfaces (interfaces)

3.2 Parameter and property types

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

3.2.1 Referenced YAML Types

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

The following table declares the valid YAML type URIs and aliases that SHALL be used when possible 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
null	tag:yaml.org,2002:null

700 **3.2.1.1 Notes**

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

704 **3.2.2 TOSCA version**

705 TOSCA supports the concept of “reuse” of type definitions, as well as template definitions which could be version

706 and change over time. It is important to provide a reliable, normative means to represent a version string which

707 enables the comparison and management of types and templates over time. Therefore, the TOSCA TC intends to

708 provide a normative version type (string) for this purpose in future Working Drafts of this specification.

Shorthand Name	version
Type Qualified Name	tosca:version

709 **3.2.2.1 Grammar**

710 TOSCA version strings have the following grammar:

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

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

- 712 • **major_version**: is a required integer value greater than or equal to 0 (zero)
- 713 • **minor_version**: is a required integer value greater than or equal to 0 (zero).
- 714 • **fix_version**: is an optional integer value greater than or equal to 0 (zero).
- 715 • **qualifier**: is an optional string that indicates a named, pre-release version of the associated code that
- 716 has been derived from the version of the code identified by the combination **major_version**,
- 717 **minor_version** and **fix_version** numbers.
- 718 • **build_version**: is an optional integer value greater than or equal to 0 (zero) that can be used to further
- 719 qualify different build versions of the code that has the same **qualifer_string**.

720 **3.2.2.2 Version Comparison**

- 721 • When comparing TOSCA versions, all component versions (i.e., major, minor and fix) are compared in
- 722 sequence from left to right.
- 723 • TOSCA versions that include the optional qualifier are considered older than those without a qualifier.
- 724 • TOSCA versions with the same major, minor, and fix versions and have the same qualifier string, but with
- 725 different build versions can be compared based upon the build version.
- 726 • Qualifier strings are considered domain-specific. Therefore, this specification makes no recommendation
- 727 on how to compare TOSCA versions with the same major, minor and fix versions, but with different
- 728 qualifiers strings and simply considers them different named branches derived from the same code.

729 **3.2.2.3 Examples**

730 Example of a version with

```
# 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
```

3.2.2.4 Notes

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

3.2.2.5 Additional Requirements

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

3.2.3 TOCSA range type

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

Shorthand Name	range
Type Qualified Name	tosca:range

3.2.3.1 Grammar

TOSCA range values have the following grammar:

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

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

- **lower_bound**: is a required integer value that denotes the lower boundary of the range.
- **upper_bound**: is a required integer value that denotes the upper boundary of the range. This value must be greater than **lower_bound**.

3.2.3.2 Keywords:

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.

3.2.3.3 Examples

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 ]
```

749 **3.2.4 TOSCA list type**

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

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

Shorthand Name	list
Type Qualified Name	tosca:list

755 **3.2.4.1 Grammar**

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

757 **3.2.4.1.1 Square bracket notation**

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

758 **3.2.4.1.2 Bulleted (sequenced) list notation**

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

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

- 760
 - **<list_entry_*>**: represents one entry of the list.

761 **3.2.4.2 Declaration Examples**

762 **3.2.4.2.1 List declaration using a simple type**

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

```
<some_entity>:  
  ...  
  properties:  
    listen_ports:  
      type: list  
      entry_schema:  
        description: listen port entry (simple integer type)  
        type: integer  
        constraints:  
          - max_length: 128
```

765 **3.2.4.2.2 List declaration using a complex type**

766 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

```

3.2.4.3 Definition Examples

These examples show two notation options for defining lists:

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

3.2.4.3.1 Square bracket notation

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

3.2.4.3.2 Bulleted list notation

```

listen_ports:
  - 80
  - 8080

```

3.2.5 TOSCA map type

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

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

Shorthand Name	map
Type Qualified Name	tosca:map

3.2.5.1 Grammar

TOSCA maps are normal YAML dictionaries with following grammar:

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>

```

784 3.2.5.1.2 Multi-line grammar

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

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

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

788 3.2.5.2 Declaration Examples

789 3.2.5.2.1 Map declaration using a simple type

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

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

792 3.2.5.2.2 Map declaration using a complex type

793 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
```

794 3.2.5.3 Definition Examples

795 These examples show two notation options for defining maps:

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

799 **3.2.5.3.1 Single-line notation**

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

800 **3.2.5.3.2 Multi-line notation**

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

801 **3.2.6 TOCSA scalar-unit type**

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

804 **3.2.6.1 Grammar**

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

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

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

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

809 **3.2.6.2 Additional requirements**

- 810 • **Whitespace**: any number of spaces (including zero or none) **SHALL** be allowed between the **scalar**
811 value and the **unit** value.
- 812 • It **SHALL** be considered an error if either the scalar or unit portion is missing on a property or attribute
813 declaration derived from any scalar-unit type.
- 814 • When performing constraint clause evaluation on values of the scalar-unit type, both the scalar value
815 portion and unit value portion **SHALL** be compared together (i.e., both are treated as a single value). For
816 example, if we have a property called **storage_size**. which is of type scalar-unit, a valid range constraint
817 would appear as follows:
- 818 ◦ **storage_size**: in_range [4 GB, 20 GB]

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

821 **3.2.6.3 Concrete Types**

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

822

823 The scalar-unit type grammar is abstract and has two recognized concrete types in TOSCA:

- 824 • **scalar-unit.size** – used to define properties that have scalar values measured in size units.
- 825 • **scalar-unit.time** – used to define properties that have scalar values measured in size units.

826 • **scalar-unit.frequency** – used to define properties that have scalar values measured in units per second.

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

828 3.2.6.4 scalar-unit.size

829 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)

830 3.2.6.4.2 Examples

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

831 3.2.6.4.3 Notes

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

840 3.2.6.5 scalar-unit.time

841 3.2.6.5.1 Recognized Units

Unit	Usage	Description
d	time	days
h	time	hours

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

842 3.2.6.5.2 Examples

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

843 3.2.6.5.3 Notes

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

850 3.2.6.6 scalar-unit.frequency

851 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.

852 3.2.6.6.2 Examples

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

853 3.2.6.6.3 Notes

- 854 • The value for Hertz (Hz) is the International Standard Unit (ISU) as described by the Bureau International
855 des Poids et Mesures (BIPM) in the “*SI Brochure: The International System of Units (SI) [8th edition,*
856 *2006; updated in 2014]*”, <http://www.bipm.org/en/publications/si-brochure/>

857 **3.3 Normative values**

858 **3.3.1 Node States**

859 As components (i.e., nodes) of TOSCA applications are deployed, instantiated and orchestrated over their
860 lifecycle using normative lifecycle operations (see section 5.7 for normative lifecycle definitions) it is important
861 define normative values for communicating the states of these components normatively between orchestration
862 and workflow engines and any managers of these applications.

863 The following table provides the list of recognized node states for TOSCA Simple Profile that would be set by the
864 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.

865 **3.3.2 Relationship States**

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

868 The following table provides the list of recognized relationship states for TOSCA Simple Profile that would be set
869 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.

870 **3.3.2.1 Notes**

- 871
- Additional states will be defined in future versions of the TOSCA Simple Profile in YAML specification.

872 3.3.3 Directives

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

874 3.3.4 Network Name aliases

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

Alias value	Description
PRIVATE	<p>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.</p> <p>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.</p>
PUBLIC	<p>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.</p> <p>A public network contains IP addresses and ports typically used to listen for incoming traffic to an application or service from the Internet.</p>

878 3.3.4.1 Usage

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

882 3.4 TOSCA Metamodel

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

885 3.4.1 Required Keynames

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

891 3.5 Reusable modeling definitions

892 3.5.1 Description definition

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

895 3.5.1.1 Keyname

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

description

897 **3.5.1.2 Grammar**

898 Description definitions have the following grammar:

```
description: <string>
```

899 **3.5.1.3 Examples**

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

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

902 The YAML “folded” style may also be used for multi-line descriptions which “folds” line breaks as space
903 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.
```

904 **3.5.1.4 Notes**

- 905
 - Use of “folded” style is discouraged for the YAML string type apart from when used with the **description**
- 906 keyname.

907 **3.5.2 Constraint clause**

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

911 **3.5.2.1 Operator keynames**

912 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.

Operator	Type	Value Type	Description
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.
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.

3.5.2.1.1 Comparable value types

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

3.5.2.2 Additional Requirements

- If no operator is present for a simple scalar-value on a constraint clause, it **SHALL** be interpreted as being equivalent to having the “**equal**” operator provided; however, the “**equal**” operator may be used for clarity when expressing a constraint clause.
- The “**length**” operator **SHALL** be interpreted mean “size” for set types (i.e., list, map, etc.).

3.5.2.3 Grammar

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>
```

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

- operator**: represents a required operator from the specified list shown above (see section 3.5.2.1 “Operator keynames”).
- scalar_value**, **scalar_value_***: represents a required scalar (or atomic quantity) that can hold only one value at a time. This will be a value of a primitive type, such as an integer or string that is allowed by this specification.

- 929 • **value_***: represents a required value of the operator that is not limited to scalars.
930 • **regular_expression_value**: represents a regular expression (string) value.

931 3.5.2.4 Examples

932 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
```

933 3.5.2.5 Additional Requirements

- 934 • Values provided by the operands (i.e., values and scalar values) **SHALL** be type-compatible with their
935 associated operations.
936 • Future drafts of this specification will detail the use of regular expressions and reference an appropriate
937 standardized grammar.

938 3.5.3 Property Filter definition

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

941 **3.5.3.1 Grammar**

942 Property filter definitions have one of the following grammars:

943 **3.5.3.1.1 Short notation:**

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

```
<property_name>: <property_constraint_clause>
```

945 **3.5.3.1.2 Extended notation:**

946 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>
```

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

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

953 **3.5.3.2 Additional Requirements**

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

956 **3.5.4 Node Filter definition**

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

959 **3.5.4.1 Keynames**

960 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.

961 **3.5.4.2 Additional filtering on named Capability properties**

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

3.5.4.3 Grammar

Node filter definitions have one of the following grammars:

```

<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>

```

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

- **property_filter_def_***: represents a property filter definition 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.
- **property_constraint_clause_***: represents constraint clause(s) that would be used to filter entities based upon property values.
- **capability_name_or_type_***: represents the type or name of a capability that would be used to select (filter) matching TOSCA entities based upon their existence.
- **cap_*_property_def_***: represents a property filter definition 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.

3.5.4.4 Additional requirements

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

3.5.4.5 Example

The following example is a filter that would be used to select a TOSCA [Compute](#) node based upon the values of its defined capabilities. Specifically, this filter would select Compute nodes that supported a specific range of

983 CPUs (i.e., **num_cpus** value between 1 and 4) and memory size (i.e., **mem_size** of 2 or greater) from its declared
984 “host” capability.
985

```
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 }
```

986 **3.5.5 Repository definition**

987 A repository definition defines a named external repository which contains deployment and implementation
988 artifacts that are referenced within the TOSCA Service Template.

989 **3.5.5.1 Keynames**

990 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.

991 **3.5.5.2 Grammar**

992 Repository definitions have one the following grammars:

993 **3.5.5.2.1 Single-line grammar (no credential):**

```
<repository_name>: <repository_address>
```

994 **3.5.5.2.2 Multi-line grammar**

```
<repository_name>:
  description: <repository_description>
  url: <repository_address>
  credential: <authorization_credential>
```

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

- 996
- **repository_name**: represents the required symbolic name of the repository as a [string](#).
 - 997 • **repository_description**: contains an optional description of the repository.
 - 998 • **repository_address**: represents the required URL of the repository as a string.

- 999
- 1000
- **authorization_credential**: represents the optional credentials (e.g., user ID and password) used to authorize access to the repository.

1001

3.5.5.3 Additional Requirements

- 1002
- None

1003

3.5.5.4 Example

1004

The following represents a repository definition:

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

1005

3.5.6 Artifact definition

1006

An artifact definition defines a named, typed file that can be associated with Node Type or Node Template and

1007

used by orchestration engine to facilitate deployment and implementation of interface operations.

1008

3.5.6.1 Keynames

1009

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

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

1010

3.5.6.2 Grammar

1011

Artifact definitions have one of the following grammars:

1012

3.5.6.2.1 Short notation

1013

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

1014

file URI:

```
<artifact_name>: <artifact_file_URI>
```

1015

3.5.6.2.2 Extended notation:

1016

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

1017

explicitly declared:

```
<artifact_name>:
  description: <artifact_description>
```

```
type: <artifact_type_name>
file: <artifact_file_URI>
repository: <artifact_repository_name>
deploy_path: <file_deployment_path>
```

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

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

3.5.6.3 Example

The following represents an artifact definition:

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

3.5.7 Import definition

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

3.5.7.1 Keynames

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.

3.5.7.2 Grammar

Import definitions have one the following grammars:

3.5.7.2.1 Single-line grammar:

```
<import_name>: <file_URI>
```

1039 3.5.7.2.2 Multi-line grammar

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

1040 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1041 • **import_name**: represents the required symbolic name for the imported file as a [string](#).
- 1042 • **file_uri**: contains the required name (i.e., URI) of the file to be imported as a [string](#).
- 1043 • **repository_name**: represents the optional symbolic name of the repository definition where the
- 1044 imported file can be found as a [string](#).
- 1045 • **namespace_uri**: represents the optional namespace URI to that will be applied to type definitions found
- 1046 within the imported file as a [string](#).
- 1047 • **namespace_prefix**: represents the optional namespace prefix (alias) that will be used to indicate the
- 1048 **namespace_uri** when forming a qualified name (i.e., QName) when referencing type definitions from the
- 1049 imported file as a [string](#).

1050 3.5.7.3 Additional Requirements

- 1051 • None

1052 3.5.7.4 Example

1053 The following represents how import definitions would be used for the imports keyname within a TOSCA Service
1054 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
```

1055 3.5.8 Property definition

1056 A property definition defines a named, typed value and related data that can be associated with an entity defined
1057 in this specification (e.g., Node Types, Relationship Types, Capability Types, etc.). Properties are used by
1058 template authors to provide input values to TOSCA entities which indicate their “desired state” when they are
1059 instantiated. The value of a property can be retrieved using the **get_property** function within TOSCA Service
1060 Templates.

1061 3.5.8.1.1 Attribute and Property reflection

1062 The actual state of the entity, at any point in its lifecycle once instantiated, is reflected by [Attribute definitions](#).
1063 TOSCA orchestrators automatically create an attribute for every declared property (with the same symbolic name)
1064 to allow introspection of both the desired state (property) and actual state (attribute).

1065 3.5.8.2 Keynames

1066 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 .

3.5.8.3 Status values

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

3.5.8.4 Grammar

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>

```

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

- **property_name**: represents the required symbolic name of the property as a [string](#).

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

3.5.8.5 Additional Requirements

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

3.5.8.6 Notes

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

3.5.8.7 Example

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 ]
```

3.5.9 Property assignment

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

3.5.9.1 Keynames

The TOSCA property assignment has no keynames.

1108 **3.5.9.2 Grammar**

1109 Property assignments have the following grammar:

1110 **3.5.9.2.1 Short notation:**

1111 The following single-line grammar may be used when a simple value assignment is needed:

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

1112 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1113 • **property_name**: represents the name of a property that would be used to select a property definition with the same name within on a TOSCA entity (e.g., Node Template, Relationship Template, etc.) which is declared in its declared type (e.g., a Node Type, Node Template, Capability Type, etc.).
- 1114 • **property_value, property_value_expression**: represent the type-compatible value to assign to the named property. Property values may be provided as the result from the evaluation of an expression or a function.

1119 **3.5.10 Attribute definition**

1120 An attribute definition defines a named, typed value that can be associated with an entity defined in this specification (e.g., a Node, Relationship or Capability Type). Specifically, it is used to expose the “actual state” of some property of a TOSCA entity after it has been deployed and instantiated (as set by the TOSCA orchestrator). Attribute values can be retrieved via the **get_attribute** function from the instance model and used as values to other entities within TOSCA Service Templates.

1125 **3.5.10.1 Attribute and Property reflection**

1126 TOSCA orchestrators automatically create [Attribute definitions](#) for any [Property definitions](#) declared on the same TOSCA entity (e.g., nodes, node capabilities and relationships) in order to make accessible the actual (i.e., the current state) value from the running instance of the entity.

1129 **3.5.10.2 Keynames**

1130 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 .

1131 **3.5.10.3 Grammar**

1132 Attribute definitions have the following grammar:

```
attributes:  
  <attribute_name>:
```

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

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

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

3.5.10.4 Additional Requirements

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

3.5.10.5 Notes

- 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.

3.5.10.6 Example

The following represents a required attribute definition:

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

3.5.11 Attribute assignment

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

3.5.11.1 Keynames

The TOSCA attribute assignment has no keynames.

3.5.11.2 Grammar

Attribute assignments have the following grammar:

1166 **3.5.11.2.1 Short notation:**

1167 The following single-line grammar may be used when a simple value assignment is needed:

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

1168 **3.5.11.2.2 Extended notation:**

1169 The following multi-line grammar may be used when a value assignment requires keys in addition to a simple
1170 value assignment:

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

1171 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1172 • **attribute_name**: represents the name of an attribute that would be used to select an attribute
1173 definition with the same name within on a TOSCA entity (e.g., Node Template, Relationship Template,
1174 etc.) which is declared (or reflected from a Property definition) in its declared type (e.g., a Node Type,
1175 Node Template, Capability Type, etc.).
- 1176 • **attribute_value, attribute_value_expression**: represent the type-compatible value to assign to
1177 the named attribute. Attribute values may be provided as the result from the evaluation of an expression
1178 or a function.
- 1179 • **attribute_description**: represents the optional [description](#) of the attribute.

1180 **3.5.11.3 Additional requirements**

- 1181 • Attribute values **MAY** be provided by the underlying implementation at runtime when requested by the
1182 get_attribute function or it **MAY** be provided through the evaluation of expressions and/or functions that
1183 derive the values from other TOSCA attributes (also at runtime).

1184 **3.5.12 Parameter definition**

1185 A parameter definition is essentially a TOSCA property definition; however, it also allows a value to be assigned
1186 to it (as for a TOSCA property assignment). In addition, in the case of output parameters, it can optionally inherit
1187 the data type of the value assigned to it rather than have an explicit data type defined for it.

1188 **3.5.12.1 Keynames**

1189 The TOSCA parameter definition has all the keynames of a TOSCA Property definition, but in addition includes
1190 the following additional or changed keynames:

Keyname	Require d	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.

1191 **3.5.12.2 Grammar**

1192 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>

```

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

- **parameter_name**: represents the required symbolic name of the parameter as a [string](#).
- **parameter_description**: represents the optional [description](#) of the parameter.
- **parameter_type**: represents the optional data type of the parameter. Note, this keyname is required for a TOSCA Property definition, but is not for a TOSCA Parameter definition.
- **parameter_value**, **parameter_value_expression**: represent 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.
- **parameter_required**: represents an optional [boolean](#) value (true or false) indicating whether or not the parameter is required. If this keyname is not present on a parameter definition, then the property SHALL be considered **required** (i.e., true) by **default**.
- **default_value**: contains a type-compatible value that may be used as a default if not provided by another means.
- **status_value**: a [string](#) that contains a keyword that indicates the status of the parameter relative to the specification or implementation.
- **parameter_constraints**: represents the optional [sequenced](#) list of one or more [constraint clauses](#) on the parameter definition.
- **entry_description**: represents the optional [description](#) of the entry schema.
- **entry_type**: represents the required type name for entries in a [list](#) or [map](#) parameter type.
- **entry_constraints**: represents the optional [sequenced](#) list of one or more [constraint clauses](#) on entries in a [list](#) or [map](#) parameter type.

3.5.12.3 Additional Requirements

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

3.5.12.4 Example

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 ]

```

1222 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 ] }

```

1223

1224 3.5.13 Operation definition

1225 An operation definition defines a named function or procedure that can be bound to an implementation artifact
1226 (e.g., a script).

1227 3.5.13.1 Keynames

1228 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.

1229 The following is the list of recognized keynames to be used with the **implementation** keyname within a TOSCA
1230 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).

1231 3.5.13.2 Grammar

1232 Operation definitions have the following grammars:

1233 3.5.13.2.1 Short notation

1234 The following single-line grammar may be used when only an operation's implementation artifact is needed:

```
<operation_name>: <implementation_artifact_name>
```

3.5.13.2.2 Extended notation for use in Type definitions

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

```
<operation_name>:
  description: <operation_description>
  implementation: <implementation_artifact_name>
  inputs:
    <property_definitions>
```

3.5.13.2.3 Extended notation for use in Template definitions

The following multi-line grammar may be used in Node or Relationship Template definitions when there 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>
```

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

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

3.5.13.3 Additional requirements

- The default sub-classing behavior for implementations of operations SHALL be override. That is, implementation artifacts assigned in subclasses override any defined in its parent class.

- 1261
- 1262
- 1263
- 1264
- 1265
- Template authors MAY provide property assignments on operation inputs on templates that do not necessarily have a property definition defined in its corresponding type.
 - Implementation artifact file names (e.g., script filenames) may include file directory path names that are relative to the TOSCA service template file itself when packaged within a TOSCA Cloud Service ARchive (CSAR) file.

1266

3.5.13.4 Examples

1267

3.5.13.4.1 Single-line implementation example

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

1268

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
```

1269

3.5.14 Interface definition

1270

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

1271

3.5.14.1 Keynames

1272

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.

1273

3.5.14.2 Grammar

1274

Interface definitions have the following grammar:

1275

3.5.14.2.1 Extended notation for use in Type definitions

1276

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>

```

3.5.14.2.2 Extended notation for use in Template definitions

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

```

<interface_definition_name>:
  inputs:
    <property_assignments>
    <operation_definitions>

```

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

- **interface_definition_name**: represents the required symbolic name of the interface as a [string](#).
- **interface_type_name**: represents the required name of the Interface Type for the interface definition.
- **property_definitions**: represents the optional list of [property definitions](#) (i.e., parameters) which the TOSCA orchestrator would make available (i.e., or pass) to all defined operations.
 - *This means these properties and their values would be accessible to the implementation artifacts (e.g., scripts) associated to each operation during their execution.*
- **property_assignments**: represents the optional list of [property assignments](#) for passing parameters to Node or Relationship Template operations providing values for properties defined in their respective type definitions.
- **operation_definitions**: represents the required name of one or more [operation definitions](#).

3.6 Type-specific definitions

3.6.1 Capability definition

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

3.6.1.1 Keynames

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.

Keyname	Required	Type	Constraints	Description
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.

3.6.1.2 Grammar

Capability definitions have one of the following grammars:

3.6.1.2.1 Short notation

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

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

3.6.1.2.2 Extended notation

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

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

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

- **capability_definition_name**: represents the symbolic name of the capability as a [string](#).
- **capability_type**: represents the required name of a [capability type](#) the capability definition is based upon.
- **capability_description**: represents the optional [description](#) of the capability definition.
- **property_definitions**: represents the optional list of [property definitions](#) for the capability definition.
- **attribute_definitions**: represents the optional list of [attribute definitions](#) for the capability definition.
- **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.

3.6.1.3 Examples

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

3.6.1.3.1 Simple notation example

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

1315 **3.6.1.3.2 Full notation example**

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

1316 **3.6.1.4 Additional requirements**

- 1317
- 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.
- 1318
- Capability symbolic names SHALL be unique; it is an error if a capability name is found to occur more than once.
- 1319
- 1320
- 1321

1322 **3.6.1.5 Notes**

- 1323
- The Capability Type, in this example **MyCapabilityTypeName**, would be defined elsewhere and have an integer property named **limit**.
- 1324
- This definition directly maps to the **CapabilitiesDefinition** of the Node Type entity as defined in the [TOSCA v1.0 specification](#).
- 1325
- 1326

1327 **3.6.2 Requirement definition**

1328 The Requirement definition describes a named requirement (dependencies) of a TOSCA Node Type or Node
1329 template which needs to be fulfilled by a matching Capability definition declared by another TOSCA modelable
1330 entity. The requirement definition may itself include the specific name of the fulfilling entity (explicitly) or provide
1331 an abstract type, along with additional filtering characteristics, that a TOSCA orchestrator can use to fulfil the
1332 capability at runtime (implicitly).

1333 **3.6.2.1 Keynames**

1334 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 fulfil 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 fulfil the requirement.
relationships	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.

1335 **3.6.2.1.1 Additional Keynames for multi-line relationship grammar**

1336 The Requirement definition contains the Relationship Type information needed by TOSCA Orchestrators to
1337 construct relationships to other TOSCA nodes with matching capabilities; however, it is sometimes recognized
1338 that additional properties may need to be passed to the relationship (perhaps for configuration). In these cases,
1339 additional grammar is provided so that the Node Type may declare additional Property definitions to be used as
1340 inputs to the Relationship Type's declared interfaces (or 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.

1341 **3.6.2.2 Grammar**

1342 Requirement definitions have one of the following grammars:

1343 **3.6.2.2.1 Simple grammar (Capability Type only)**

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

1344 **3.6.2.2.2 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> ]
```

1345 **3.6.2.2.3 Extended grammar for declaring Property Definitions on the relationship's**
1346 **Interfaces**

1347 The following additional multi-line grammar is provided for the relationship keyname in order to declare new
1348 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>
```

1349 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1350 • **requirement_name**: represents the required symbolic name of the requirement definition as a string.
- 1351 • **capability_type_name**: represents the required name of a Capability type that can be used to fulfil the
- 1352 requirement.

- **node_type_name**: represents the optional name of a TOSCA Node Type that contains the Capability Type definition the requirement can be fulfilled by.
- **relationship_type_name**: represents the optional name of a [Relationship Type](#) to be used to construct a relationship between this requirement definition (i.e., in the source node) to a matching capability definition (in a target node).
- **min_occurrences, max_occurrences**: represents the optional minimum and maximum occurrences of the requirement (i.e., its cardinality).
- **interface_definitions**: represents one or more already declared interface definitions in the Relationship Type (as declared on the **type** keyname) allowing for the declaration of new Property definition for these interfaces or for specific Operation definitions of these interfaces.

3.6.2.3 Additional Requirements

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

3.6.2.4 Notes

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

3.6.2.5 Requirement Type definition is a tuple

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

1. Node Type (optional)
2. Relationship Type (optional)
3. Capability Type (required)

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

3.6.2.5.1 Property filter

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

3.6.3 Artifact Type

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

3.6.3.1 Keynames

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

Keyname	Required	Type	Description
derived_from	no	string	An optional parent Artifact Type name the Artifact Type derives from.
version	no	version	An optional version for the Artifact Type definition.
description	no	description	An optional description for the Artifact Type.
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.

3.6.3.2 Grammar

Artifact Types have following grammar:

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

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

- **artifact_type_name**: represents the name of the Artifact Type being declared as a [string](#).
- **parent_artifact_type_name**: represents the [name](#) of the [Artifact Type](#) this Artifact Type definition derives from (i.e., its “parent” type).
- **version_number**: represents the optional TOSCA [version](#) number for the Artifact Type.
- **artifact_description**: represents the optional [description](#) string for the Artifact Type.
- **mime_type_string**: represents the optional Multipurpose Internet Mail Extensions (MIME) standard string value that describes the file contents for this type of Artifact Type as a [string](#).
- **file_extensions**: represents the optional list of one or more recognized file extensions for this type of artifact type as [strings](#).
- **property_definitions**: represents the optional list of [property definitions](#) for the artifact type.

3.6.3.3 Examples

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

3.6.4 Interface Type

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

1411 **3.6.4.1 Keynames**

1412 The following is the list of recognized keynames for a TOSCA Interface Type definition:

Keyname	Required	Type	Description
version	no	version	An optional version for the Interface Type definition.
description	no	description	An optional description for the Interface Type.
inputs	no	list of property definitions	The optional list of input parameter definitions.

1413 **3.6.4.2 Grammar**

1414 Interface Types have following grammar:

```
<interface_type_name>:
  version: <version_number>
  description: <interface_description>
  inputs:
    <property_definitions>
    <operation_definitions>
```

1415 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1416
- **interface_type_name**: represents the required name of the interface as a [string](#).
 - 1417 • **version_number**: represents the optional TOSCA [version](#) number for the Interface Type.
 - 1418 • **interface_description**: represents the optional [description](#) string for the Interface Type.
 - 1419 • **property_definitions**: represents the optional list of [property definitions](#) (i.e., parameters) which the
 - 1420 TOSCA orchestrator would make available (i.e., or pass) to all implementation artifacts for operations
 - 1421 declared on the interface during their execution.
 - 1422 • **operation_definitions**: represents the required list of one or more [operation definitions](#).

1423 **3.6.4.3 Example**

1424 The following example shows a custom interface used to define multiple configure operations.

```
mycompany.mytypes.myinterfaces.MyConfigure:
  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
```

1425 **3.6.4.4 Additional Requirements**

- 1426
- Interface Types **MUST NOT** include any implementations for defined operations; that is, the implementation keyname is invalid.
 - 1427
 - 1428 • The **inputs** keyname is reserved and **SHALL NOT** be used for an operation name.

1429 **3.6.4.5 Notes**

- 1430 • The TOSCA Simple Profile specification does not yet provide a means to derive or extend an Interface
1431 Type from another Interface Type.

1432 **3.6.5 Data Type**

1433 A Data Type definition defines the schema for new named datatypes in TOSCA.

1434 **3.6.5.1 Keynames**

1435 The following is the list of recognized keynames for a TOSCA Data Type definition:

Keyname	Required	Type	Description
derived_from	no	string	The optional key used when a datatype is derived from an existing TOSCA Data Type.
version	no	version	An optional version for the Data Type definition.
description	no	description	The optional description for the Data Type.
constraints	no	list of constraint clauses	The optional list of <i>sequenced</i> 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.

1436 **3.6.5.2 Grammar**

1437 Data Types have the following grammar:

```
<data_type_name>:  
  derived_from: <existing_type_name>  
  version: <version_number>  
  description: <datatype_description>  
  constraints:  
    - <type_constraints>  
  properties:  
    <property_definitions>
```

1438 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1439 • **data_type_name**: represents the required symbolic name of the Data Type as a **string**.
1440 • **version_number**: represents the optional TOSCA **version** number for the Data Type.
1441 • **datatype_description**: represents the optional **description** for the Data Type.
1442 • **existing_type_name**: represents the optional name of a valid TOSCA type this new Data Type would
1443 derive from.
1444 • **type_constraints**: represents the optional *sequenced* list of one or more type-compatible **constraint**
1445 **clauses** that restrict the Data Type.
1446 • **property_definitions**: represents the optional list of one or more **property definitions** that provide the
1447 schema for the Data Type.

1448 **3.6.5.3 Additional Requirements**

- 1449 • A valid datatype definition **MUST** have either a valid **derived_from** declaration or at least one valid
1450 property definition.

- 1451
- 1452
- 1453
- Any **constraint** clauses **SHALL** be type-compatible with the type declared by the **derived_from** keyname.
 - If a **properties** keyname is provided, it **SHALL** contain one or more valid property definitions.

1454

3.6.5.4 Examples

1455

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

1456

3.6.5.4.1 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
```

1457

3.6.5.4.2 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
```

1458

3.6.6 Capability Type

1459

A Capability Type is a reusable entity that describes a kind of capability that a Node Type can declare to expose.

1460

Requirements (implicit or explicit) that are declared as part of one node can be matched to (i.e., fulfilled by) the

1461

Capabilities declared by another node.

1462

3.6.6.1 Keynames

1463

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

Keyname	Required	Type	Description
derived_from	no	string	An optional parent capability type name this new Capability Type derives from.
version	no	version	An optional version for the Capability Type definition.
description	no	description	An optional description for the Capability Type.

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.

3.6.6.2 Grammar

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> ]

```

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

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

3.6.6.3 Example

```

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

```

1479 **3.6.7 Requirement Type**

1480 A Requirement Type is a reusable entity that describes a kind of requirement that a Node Type can declare to
1481 expose. The TOSCA Simple Profile seeks to simplify the need for declaring specific Requirement Types from
1482 nodes and instead rely upon nodes declaring their features sets using TOSCA Capability Types along with a
1483 named Feature notation.

1484 Currently, there are no use cases in this TOSCA Simple Profile in YAML specification that utilize an independently
1485 defined Requirement Type. This is a desired effect as part of the simplification of the TOSCA v1.0 specification.

1486 **3.6.8 Node Type**

1487 A Node Type is a reusable entity that defines the type of one or more Node Templates. As such, a Node Type
1488 defines the structure of observable properties via a *Properties Definition, the Requirements and Capabilities of the*
1489 *node as well as its supported interfaces.*

1490 **3.6.8.1 Keynames**

1491 The following is the list of recognized keynames for a TOSCA Node Type definition:

Keyname	Required	Definition/Type	Description
derived_from	no	string	An optional parent Node Type name this new Node Type derives from.
version	no	version	An optional version for the Node Type definition.
description	no	description	An optional description for the Node Type.
properties	no	list of property definitions	An optional list of property definitions for the Node Type.
attributes	no	list of attribute definitions	An optional list of attribute 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.

1492 **3.6.8.2 Grammar**

1493 Node Types have following grammar:

```
<node_type_name>:  
  derived_from: <parent_node_type_name>  
  version: <version_number>  
  description: <node_type_description>  
  properties:  
    <property_definitions>  
  attributes:  
    <attribute_definitions>  
  requirements:  
    - <requirement_definitions>
```



```

capabilities:
  <capability_definitions>
interfaces:
  <interface_definitions>
artifacts:
  <artifact_definitions>

```

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

- **node_type_name**: represents the required symbolic name of the Node Type being declared.
- **parent_node_type_name**: represents the name ([string](#)) of the [Node Type](#) this Node Type definition derives from (i.e., its “parent” type).
- **version_number**: represents the optional TOSCA [version](#) number for the Node Type.
- **node_type_description**: represents the optional [description](#) string for the corresponding **node_type_name**.
- **property_definitions**: represents the optional list of [property definitions](#) for the Node Type.
- **attribute_definitions**: represents the optional list of [attribute definitions](#) for the Node Type.
- **requirement_definitions**: represents the optional [sequenced](#) list of [requirement definitions](#) for the Node Type.
- **capability_definitions**: represents the optional list of [capability definitions](#) for the Node Type.
- **interface_definitions**: represents the optional list of one or more [interface definitions](#) supported by the Node Type.
- **artifact_definitions**: represents the optional list of [artifact definitions](#) for the Node Type.

3.6.8.3 Additional Requirements

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

3.6.8.4 Best Practices

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

3.6.8.5 Example

```

my_company.my_types.my_app_node_type:
  derived_from: tosca.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

```

3.6.9 Relationship Type

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

3.6.9.1 Keynames

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

Keyname	Required	Definition/Type	Description
derived_from	no	string	An optional parent Relationship Type name the Relationship Type derives from.
version	no	version	An optional version for the Relationship Type definition.
description	no	description	An optional description for the Relationship Type.
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.

3.6.9.2 Grammar

Relationship Types have following grammar:

```

<relationship_type_name>:
  derived_from: <parent_relationship_type_name>
  version: <version_number>
  description: <relationship_description>
  properties:
    <property_definitions>
  attributes:
    <attribute_definitions>
  interfaces:

```

```

<interface_definitions>
valid_target_types: [ <capability_type_names> ]

```

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

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

3.6.9.3 Best Practices

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

3.6.9.4 Examples

```

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

```

3.6.10 Group Type

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

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

3.6.10.1 Keynames

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

Keyname	Required	Type	Description
derived_from	no	string	An optional parent Group Type name the Group Type derives from.

Keyname	Required	Type	Description
version	no	version	An optional version for the Group Type definition.
description	no	description	The optional description for the Group Type.
properties	no	list of property definitions	An optional list of property definitions for the Group Type.
targets	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 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 "InvitesTo" (group).
interfaces	no	list of interface definitions	An optional list of interface definitions supported by the Group Type.

3.6.10.2 Grammar

Group Types have one the following grammars:

```

<group_type_name>:
  derived_from: <parent_group_type_name>
  version: <version_number>
  description: <group_description>
  properties:
    <property_definitions>
  targets: [ <list_of_valid_target_types> ]
  interfaces:
    <interface_definitions>

```

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

- **group_type_name**: represents the required symbolic name of the Group Type being declared as a [string](#).
- **parent_group_type_name**: represents the name ([string](#)) of the [Group Type](#) this Group Type definition derives from (i.e., its "parent" type).
- **version_number**: represents the optional TOSCA [version](#) number for the Group Type.
- **group_description**: represents the optional description string for the corresponding **group_type_name**.
- **property_definitions**: represents the optional list of [property definitions](#) for the Group Type.
- **list_of_valid_target_types**: represents the optional list of TOSCA types (i.e., Node or Capability Types) that are valid target types for being added to (i.e., members of) the Group Type.
- **interface_definitions**: represents the optional list of one or more [interface definitions](#) supported by the Group Type.

3.6.10.3 Additional Requirements

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

1583 **3.6.10.4 Example**

1584 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
    targets: [ tosca.nodes.Compute ]
```

1585 **3.6.11 Policy Type**

1586 A Policy Type defines a type of requirement that affects or governs an application or service's topology at some
1587 stage of its lifecycle, but is not explicitly part of the topology itself (i.e., it does not prevent the application or
1588 service from being deployed or run if it did not exist).

1589 **3.6.11.1 Keynames**

1590 The following is the list of recognized keynames for a TOSCA Policy Type definition:

Keyname	Required	Type	Description
derived_from	no	string	An optional parent Policy Type name the Policy Type derives from.
version	no	version	An optional version for the Policy Type definition.
description	no	description	The optional description for the Policy Type.
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).

1591 **3.6.11.2 Grammar**

1592 Policy Types have one the following grammars:

```
<policy_type_name>:
  derived_from: <parent_policy_type_name>
  version: <version_number>
  description: <policy_description>
  properties:
    <property_definitions>
  targets: [ <list_of_valid_target_types> ]
```

1593 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1594
- **policy_type_name**: represents the required symbolic name of the Policy Type being declared as a string.
 - **parent_policy_type_name**: represents the name (string) of the Policy Type this Policy Type definition derives from (i.e., its "parent" type).
 - **version_number**: represents the optional TOSCA version number for the Policy Type.
- 1595
- 1596
- 1597
- 1598

- 1599
- 1600
- 1601
- 1602
- 1603
- **policy_description**: represents the optional description string for the corresponding **policy_type_name**.
 - **property_definitions**: represents the optional list of [property definitions](#) for the Policy Type.
 - **list_of_valid_target_types**: represents the optional list of TOSCA types (i.e., Group or Node Types) that are valid targets for this Policy Type.

1604

3.6.11.3 Additional Requirements

- 1605
- None

1606

3.6.11.4 Example

1607

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: tosca.policies.Root
```

1608

3.7 Template-specific definitions

1609

1610

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

1611

3.7.1 Capability assignment

1612

1613

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

1614

3.7.1.1 Keynames

1615

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.

1616

3.7.1.2 Grammar

1617

Capability assignments have one of the following grammars:

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

1618

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

- 1619
- 1620
- **capability_definition_name**: represents the symbolic name of the capability as a [string](#).
 - **property_assignments**: represents the optional list of [property assignments](#) for the capability definition.

1621 • **attribute_assignments**: represents the optional list of **attribute assignments** for the capability definition.

1622 **3.7.1.3 Example**

1623 The following example shows a capability assignment:

1624 **3.7.1.3.1 Notation example**

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

1625 **3.7.2 Requirement assignment**

1626 A Requirement assignment allows template authors to provide either concrete names of TOSCA templates or
1627 provide abstract selection criteria for providers to use to find matching TOSCA templates that are used to fulfill a
1628 named requirement's declared TOSCA Node Type.

1629 **3.7.2.1 Keynames**

1630 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 fulfil the target node requirement.• Node Type name that the provider will use to select a type-compatible node template to fulfil 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.

1631 The following is the list of recognized keynames for a TOSCA requirement assignment's **relationship** keyname
1632 which is used when Property assignments need to be provided to inputs of declared interfaces 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.

Keyname	Required	Type	Description
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.

3.7.2.2 Grammar

Named requirement assignments have one of the following grammars:

3.7.2.2.1 Short notation:

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

```
<requirement_name>: <node_template_name>
```

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

3.7.2.2.2 Extended notation:

The following grammar would be used if the requirement assignment needs to provide more information 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 ]
```

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

The following additional multi-line grammar is provided for the relationship keyname in order to provide 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>
```

Examples of uses for the extended requirement assignment grammar include:

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

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

- **requirement_name**: represents the symbolic name of a requirement assignment as a [string](#).
- **node_template_name**: represents the optional name of a Node Template that contains the capability this requirement will be fulfilled by.
- **relationship_template_name**: represents the optional name of a [Relationship Type](#) to be used when relating the requirement appears to the capability in the target node.
- **capability_symbolic_name**: represents the optional ordered list of specific, required capability type or named capability definition within the target Node Type or Template.
- **node_type_name**: represents the optional name of a TOSCA Node Type the associated named requirement can be fulfilled by. This must be a type that is compatible with the Node Type declared on the matching requirement (same symbolic name) the requirement’s Node Template is based upon.
- **relationship_type_name**: represents the optional name of a [Relationship Type](#) that is compatible with the Capability Type in the target node.
- **property_assignments**: represents the optional list of property value assignments for the declared relationship.
- **interface_assignments**: represents the optional list of interface definitions for the declared relationship used to provide property assignments on inputs of interfaces and operations.
- **capability_type_name**: represents the optional name of a Capability Type definition within the target Node Type this requirement needs to form a relationship with.
- **node_filter_definition**: represents the optional [node filter](#) TOSCA orchestrators would use to fulfill the requirement for selecting a target node. Note that this SHALL only be valid if the **node** keyname’s value is a Node Type and is invalid if it is a Node Template.

3.7.2.3 Examples

3.7.2.3.1 Example 1 – Abstract hosting requirement on a Node Type

A web application node template named ‘`my_application_node_template`’ of type `WebApplication` declares a requirement named ‘`host`’ that needs to be fulfilled by any node that derives from the node 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
```

1683 In this case, the node template's type is **WebApplication** which already declares the Relationship Type
1684 **HostedOn** to use to relate to the target node and the Capability Type of **Container** to be the specific target of the
1685 requirement in the target node.

1686 3.7.2.3.2 Example 2 - Requirement with Node Template and a custom Relationship Type

1687 This example is similar to the previous example; however, the requirement named 'database' describes a
1688 requirement for a connection to a database endpoint (**Endpoint.Database**) Capability Type in a named node
1689 template (**my_database**). However, the connection requires a custom Relationship Type
1690 (**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
```

1691 3.7.2.3.3 Example 3 - Requirement for a Compute node with additional selection criteria 1692 (filter)

1693 This example shows how to extend an abstract 'host' requirement for a **Compute** node with a
1694 filter definition that further constrains TOSCA orchestrators to include additional
1695 properties and capabilities on the target node when fulfilling the 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 ] }
```

1696 **3.7.3 Node Template**

1697 A Node Template specifies the occurrence of a manageable software component as part of an application's
1698 topology model which is defined in a TOSCA Service Template. A Node template is an instance of a specified
1699 Node Type and can provide customized properties, constraints or operations which override the defaults provided
1700 by its Node Type and its implementations.

1701 **3.7.3.1 Keynames**

1702 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.
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 <i>sequenced</i> list of requirement assignments for the Node Template.
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.

1703 **3.7.3.2 Grammar**

```
<node_template_name>:
  type: <node_type_name>
  description: <node_template_description>
  directives: [<directives>]
  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>

```

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

- **node_template_name**: represents the required symbolic name of the Node Template being declared.
- **node_type_name**: represents the name of the Node Type the Node Template is based upon.
- **node_template_description**: represents the optional [description](#) string for Node Template.
- **directives**: represents the optional list of processing instruction keywords (as strings) for use by tooling and orchestrators.
- **property_assignments**: represents the optional list of [property assignments](#) for the Node Template that provide values for properties defined in its declared Node Type.
- **attribute_assignments**: represents the optional list of [attribute assignments](#) for the Node Template that provide values for attributes defined in its declared Node Type.
- **requirement_assignments**: represents the optional [sequenced](#) list of [requirement assignments](#) for the Node Template that allow assignment of type-compatible capabilities, target nodes, relationships and target (node filters) for use when fulfilling the requirement at runtime.
- **capability_assignments**: represents the optional list of [capability assignments](#) for the Node Template that augment those provided by its declared Node Type.
- **interface_definitions**: represents the optional list of [interface definitions](#) for the Node Template that [augment](#) those provided by its declared Node Type.
- **artifact_definitions**: represents the optional list of [artifact definitions](#) for the Node Template that augment those provided by its declared Node Type.
- **node_filter_definition**: represents the optional [node filter](#) TOSCA orchestrators would use for selecting a matching node template.
- **source_node_template_name**: represents the optional (symbolic) name of another node template to copy into (all keynames and values) and use as a basis for this node template.

3.7.3.3 Additional requirements

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

3.7.3.4 Example

```

node_templates:
  mysql:
    type: tosa.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

```

3.7.4 Relationship Template

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

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.
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.

3.7.4.1 Grammar

```

<relationship_template_name>:
  type: <relationship_type_name>
  description: <relationship_type_description>
  properties:
    <property_assignments>
  attributes:
    <attribute_assignments>
  interfaces:
    <interface_definitions>
  copy:
    <source_relationship_template_name>

```

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

- relationship_template_name**: represents the required symbolic name of the Relationship Template being declared.
- relationship_type_name**: represents the name of the Relationship Type the Relationship Template is based upon.

- **relationship_template_description**: represents the optional [description](#) string for the Relationship Template.
- **property_assignments**: represents the optional list of [property assignments](#) for the Relationship Template that provide values for properties defined in its declared Relationship Type.
- **attribute_assignments**: represents the optional list of [attribute assignments](#) for the Relationship Template that provide values for attributes defined in its declared Relationship Type.
- **interface_definitions**: represents the optional list of [interface definitions](#) for the Relationship Template that augment those provided by its declared Relationship Type.
- **source_relationship_template_name**: represents the optional (symbolic) name of another relationship template to copy into (all keynames and values) and use as a basis for this relationship template.

3.7.4.2 Additional requirements

- 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).

3.7.4.3 Example

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

3.7.5 Group definition

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

3.7.5.1 Keynames

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.
properties	no	list of property assignments	An optional list of property value assignments for the group definition.
targets	yes	list of string	The required 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.

3.7.5.2 Grammar

Group definitions have one the following grammars:

```
<group_name>:
  type: <group_type_name>
```

```

description: <group_description>
properties:
  <property_assignments>
targets: [ <list_of_node_templates> ]
interfaces:
  <interface_definitions>

```

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

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

3.7.5.3 Additional Requirements

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

3.7.5.4 Example

The following represents a group definition:

```

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

```

3.7.6 Policy definition

A policy definition defines a policy that can be associated with a TOSCA topology or top-level entity definition (e.g., group definition, node template, etc.).

3.7.6.1 Keynames

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.
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.

1789 **3.7.6.2 Grammar**

1790 Policy definitions have one the following grammars:

```
<policy_name>:
  type: <policy_type_name>
  description: <policy_description>
  properties:
    <property_assignments>
  targets: [<list_of_policy_targets>]
```

1791 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1792
- 1793 • **policy_name**: represents the required symbolic name of the policy as a [string](#).
 - 1794 • **policy_type_name**: represents the name of the policy the definition is based upon.
 - 1795 • **policy_description**: contains an optional description of the policy.
 - 1796 • **property_assignments**: represents the optional list of [property assignments](#) for the policy definition
 - 1797 that provide values for properties defined in its declared Policy Type.
 - 1798 • **list_of_policy_targets**: represents the optional list of names of node templates or groups that the

1799 **3.7.6.3 Additional Requirements**

- 1800
- None

1801 **3.7.6.4 Example**

1802 The following represents a policy definition:

```
policies:
- my_compute_placement_policy:
  type: tosa.policies.placement
  description: Apply my placement policy to my application's servers
  targets: [ my_server_1, my_server_2 ]
```

1803 **3.8 Topology Template definition**

1804 This section defines the topology template of a cloud application. The main ingredients of the topology template
1805 are node templates representing components of the application and relationship templates representing links
1806 between the components. These elements are defined in the nested **node_templates** section and the nested
1807 **relationship_templates** sections, respectively. Furthermore, a topology template allows for defining input
1808 parameters, output parameters as well as grouping of node templates.

1809 **3.8.1 Keynames**

1810 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.

Keyname	Required	Type	Description
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.
substitution_mappings	no	N/A	<p>An optional declaration that exports the topology template as an implementation of a Node type.</p> <p>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.</p>

3.8.2 Grammar

The overall grammar of the **topology_template** section is shown below.—Detailed grammar definitions 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>

  # 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>

```

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

- **template_description**: represents the optional [description](#) string for Topology Template.

- **input_parameter_list**: represents the optional list of input parameters (i.e., as property definitions) for the Topology Template.
- **output_parameter_list**: represents the optional list of output parameters (i.e., as property definitions) for the Topology Template.
- **group_definition_list**: represents the optional list of [group definitions](#) whose members are node templates that also are defined within this Topology Template.
- **policy_definition_list**: represents the optional sequenced list of policy definitions for the Topology Template.
- **node_template_list**: represents the optional list of [node template](#) definitions for the Topology Template.
- **relationship_template_list**: represents the optional list of [relationship templates](#) for the Topology Template.
- **node_type_name**: represents the optional name of a [Node Type](#) that the Topology Template implements as part of the **substitution_mappings**.
- **map_of_capability_mappings_to_expose**: represents the mappings that expose internal capabilities from node templates (within the topology template) as capabilities of the Node Type definition that is declared as part of the **substitution_mappings**.
- **map_of_requirement_mappings_to_expose**: represents the mappings of link requirements of the Node Type definition that is declared as part of the **substitution_mappings** to internal requirements implementations within node templates (declared within the topology template).

More detailed explanations for each of the Topology Template grammar's keynames appears in the sections below.

3.8.2.1 inputs

The **inputs** section provides a means to define parameters using TOSCA parameter definitions, their allowed values via constraints and default values within a TOSCA Simple Profile template. Input parameters defined in the **inputs** section of a topology template can be mapped to properties of node templates or relationship templates within the same topology template and can thus be used for parameterizing the instantiation of the topology template.

This section defines topology template-level input parameter section.

- Inputs here would ideally be mapped to BoundaryDefinitions in TOSCA v1.0.
- Treat input parameters as fixed global variables (not settable within template)
- If not in input take default (nodes use default)

3.8.2.1.1 Grammar

The grammar of the **inputs** section is as follows:

```
inputs:
  <parameter_definition_list>
```

3.8.2.1.2 Examples

This section provides a set of examples for the single elements of a topology template.

Simple **inputs** example without any constraints:

```
inputs:
  fooName:
```

```
type: string
description: Simple string typed property definition with no constraints.
default: bar
```

1855 Example of **inputs** with constraints:

```
inputs:
  SiteName:
    type: string
    description: string typed property definition with constraints
    default: My Site
    constraints:
      - min_length: 9
```

1856 3.8.2.2 node_templates

1857 The **node_templates** section lists the Node Templates that describe the (software) components that are used to
1858 compose cloud applications.

1859 3.8.2.2.1 grammar

1860 The grammar of the **node_templates** section is as follows:

```
node_templates:
  <node_template_defn_1>
  ...
  <node_template_defn_n>
```

1861 3.8.2.2.2 Example

1862 Example of **node_templates** section:

```
node_templates:
  my_webapp_node_template:
    type: WebApplication

  my_database_node_template:
    type: Database
```

1863 3.8.2.3 relationship_templates

1864 The **relationship_templates** section lists the Relationship Templates that describe the relations between
1865 components that are used to compose cloud applications.

1866

1867 Note that in the TOSCA Simple Profile, the explicit definition of relationship templates as it was required in
1868 TOSCA v1.0 is optional, since relationships between nodes get implicitly defined by referencing other node
1869 templates in the requirements sections of node templates.

1870 3.8.2.3.1 Grammar

1871 The grammar of the **relationship_templates** section is as follows:

```
relationship_templates:
  <relationship_template_defn_1>
  ...
  <relationship_template_defn_n>
```

3.8.2.3.2 Example

Example of **relationship_templates** section:

```
relationship_templates:
  my_connectsto_relationship:
    type: toska.relationships.ConnectsTo
    interfaces:
      Configure:
        inputs:
          speed: { get_attribute: [ SOURCE, connect_speed ] }
```

3.8.2.4 outputs

The **outputs** section provides a means to define the output parameters that are available from a TOSCA Simple Profile service template. It allows for exposing attributes of node templates or relationship templates within the containing **topology_template** to users of a service.

3.8.2.4.1 Grammar

The grammar of the **outputs** section is as follows:

```
outputs:
  <parameter_def_list>
```

3.8.2.4.2 Example

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 ] }
```

3.8.2.5 groups

The **groups** section allows for grouping one or more node templates within a TOSCA Service Template and for assigning special attributes like policies to the group.

3.8.2.5.1 Grammar

The grammar of the **groups** section is as follows:

```
groups:
  <group_defn_1>
  ...
```

<group_defn_n>

3.8.2.5.2 Example

The following example shows the definition of three Compute nodes in the **node_templates** section of a **topology_template** as well as the grouping of two of the Compute nodes in a group **server_group_1**.

```
node_templates:
  server1:
    type: tosa.nodes.Compute
    # more details ...

  server2:
    type: tosa.nodes.Compute
    # more details ...

  server3:
    type: tosa.nodes.Compute
    # more details ...

groups:
  # server2 and server3 are part of the same group
  - server_group_1:
      type: tosa.groups.Root
      targets: [ server2, server3 ]
```

3.8.2.6 policies

The **policies** section allows for declaring policies that can be applied to entities in the topology template.

3.8.2.6.1 Grammar

The grammar of the **policies** section is as follows:

```
policies:
  - <policy_defn_1>
  - ...
  - <policy_defn_n>
```

3.8.2.6.2 Example

The following example shows the definition of a placement policy.

```
policies:
  - my_placement_policy:
      type: mycompany.mytypes.policy.placement
```

3.8.2.7 Notes

- The parameters (properties) that are listed as part of the **inputs** block can be mapped to **PropertyMappings** provided as part of **BoundaryDefinitions** as described by the TOSCA v1.0 specification.
- The node templates listed as part of the **node_templates** block can be mapped to the list of **NodeTemplate** definitions provided as part of **TopologyTemplate** of a **ServiceTemplate** as described by the TOSCA v1.0 specification.
- The relationship templates listed as part of the **relationship_templates** block can be mapped to the list of **RelationshipTemplate** definitions provided as part of **TopologyTemplate** of a **ServiceTemplate** as described by the TOSCA v1.0 specification.
- The output parameters that are listed as part of the **outputs** section of a topology template can be mapped to **PropertyMappings** provided as part of **BoundaryDefinitions** as described by the TOSCA v1.0 specification.
 - Note, however, that TOSCA v1.0 does not define a direction (input vs. output) for those mappings, i.e. TOSCA v1.0 **PropertyMappings** are underspecified in that respect and TOSCA Simple Profile's **inputs** and **outputs** provide a more concrete definition of input and output parameters.

3.9 Service Template definition

A TOSCA Service Template (YAML) document contains element definitions of building blocks for cloud application, or complete models of cloud applications. This section describes the top-level structural elements (TOSCA keynames) along with their grammars, which are allowed to appear in a TOSCA Service Template document.

3.9.1 Keynames

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.

Keyname	Required	Type	Description
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.

3.9.1.1 Metadata keynames

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.

3.9.2 Grammar

The overall structure of a TOSCA Service Template and its top-level key collations using the TOSCA Simple Profile is shown below:

```

tosca_definitions_version: # Required TOSCA Definitions version string

# Optional metadata keyname: value pairs
metadata:
  template_name:           # Optional name of this service template
  template_author:         # Optional author of this service template
  template_version:        # 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

```

3.9.2.1 Notes

- TOSCA Service Templates do not have to contain a topology_template and MAY contain simply type definitions (e.g., Artifact, Interface, Capability, Node, Relationship Types, etc.) and be imported for use as type definitions in other TOSCA Service Templates.

3.9.3 Top-level keyname definitions

3.9.3.1 **tosca_definitions_version**

This required element provides a means to include a reference to the TOSCA Simple Profile specification within the TOSCA Definitions YAML file. It is an indicator for the version of the TOSCA grammar that should be used to parse the remainder of the document.

1934 3.9.3.1.1 Keyname

```
tosca_definitions_version
```

1935 3.9.3.1.2 Grammar

1936 Single-line form:

```
tosca_definitions_version: <tosca_simple_profile_version>
```

1937 3.9.3.1.3 Examples:

1938 TOSCA Simple Profile version 1.0 specification using the defined namespace alias (see Section 3.1):

```
tosca_definitions_version: toska_simple_yaml_1_0
```

1939 TOSCA Simple Profile version 1.0 specification using the fully defined (target) namespace (see Section 3.1):

```
tosca_definitions_version: http://docs.oasis-open.org/tosca/ns/simple/yaml/1.0
```

1940 3.9.3.2 metadata

1941 This keyname is used to associate domain-specific metadata with the Service Template. The metadata keyname
1942 allows a declaration of a map of keynames with string values.

1943 3.9.3.2.1 Keyname

```
metadata
```

1944 3.9.3.2.2 Grammar

```
metadata:  
  <map_of_string_values>
```

1945 3.9.3.2.3 Example

```
metadata:  
  creation_date: 2015-04-14  
  date_updated: 2015-05-01  
  status: developmental
```

1946

1947 3.9.3.3 template_name

1948 This optional metadata keyname can be used to declare the name of service template as a single-line string
1949 value.

1950 3.9.3.3.1 Keyname

```
template_name
```

1951 3.9.3.3.2 Grammar

```
template_name: <name string>
```

1952 3.9.3.3.3 Example

```
template_name: My service template
```

1953 3.9.3.3.4 Notes

- 1954 • Some service templates are designed to be referenced and reused by other service templates.
1955 Therefore, in these cases, the **template_name** value SHOULD be designed to be used as a unique
1956 identifier through the use of namespacing techniques.

1957 3.9.3.4 template_author

1958 This optional metadata keyname can be used to declare the author(s) of the service template as a single-line
1959 string value.

1960 3.9.3.4.1 Keyname

```
template_author
```

1961 3.9.3.4.2 Grammar

```
template_author: <author string>
```

1962 3.9.3.4.3 Example

```
template_author: My service template
```

1963 3.9.3.5 template_version

1964 This optional metadata keyname can be used to declare a domain specific version of the service template as a
1965 single-line string value.

1966 3.9.3.5.1 Keyname

```
template_version
```

1967 3.9.3.5.2 Grammar

```
template_version: <version>
```

1968 3.9.3.5.3 Example

```
template_version: 2.0.17
```

1969 3.9.3.5.4 Notes:

- 1970 • Some service templates are designed to be referenced and reused by other service templates and have a
1971 lifecycle of their own. Therefore, in these cases, a **template_version** value SHOULD be included and

1972 used in conjunction with a unique **template_name** value to enable lifecycle management of the service
1973 template and its contents.

1974 3.9.3.6 description

1975 This optional keyname provides a means to include single or multiline descriptions within a TOSCA Simple Profile
1976 template as a scalar string value.

1977 3.9.3.6.1 Keyname

```
description
```

1978 3.9.3.7 dsl_definitions

1979 This optional keyname provides a section to define macros (e.g., YAML-style macros when using the TOSCA
1980 Simple Profile in YAML specification).

1981 3.9.3.7.1 Keyname

```
dsl_definitions
```

1982 3.9.3.7.2 Grammar

```
dsl_definitions:  
  <dsl_definition_1>  
  ...  
  <dsl_definition_n>
```

1983 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
```

1984 3.9.3.8 repositories

1985 This optional keyname provides a section to define external repositories which may contain artifacts or other
1986 TOSCA Service Templates which might be referenced or imported by the TOSCA Service Template definition.

1987 3.9.3.8.1 Keyname

```
repositories
```

1988 3.9.3.8.2 Grammar

```
repositories:
  <repository_definition_1>
  ...
  <repository_definition_n>
```

1989 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/
```

1990 3.9.3.9 imports

1991 This optional keyname provides a way to import a *block sequence* of one or more TOSCA Definitions documents.
1992 TOSCA Definitions documents can contain reusable TOSCA type definitions (e.g., Node Types, Relationship
1993 Types, Artifact Types, etc.) defined by other authors. This mechanism provides an effective way for companies
1994 and organizations to define normative types and/or describe their software applications for reuse in other TOSCA
1995 Service Templates.

1996 3.9.3.9.1 Keyname

```
imports
```

1997 3.9.3.9.2 Grammar

```
imports:
  - <import_definition_1>
  - ...
  - <import_definition_n>
```

1998 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
  - more_definitions:
      file: my_defns/my_typesdefs_n.yaml
      repository: my_company_repo
      namespace_uri: http://mycompany.com/ns/tosca/2.0
      namespace_prefix: mycompany
```

1999 **3.9.3.10 artifact_types**

2000 This optional keyname lists the Artifact Types that are defined by this Service Template.

2001 **3.9.3.10.1 Keyname**

```
artifact_types
```

2002 **3.9.3.10.2 Grammar**

```
artifact_types:
  <artifact_type_defn_1>
  ...
  <artifact_type_defn_n>
```

2003 **3.9.3.10.3 Example**

```
artifact_types:
  mycompany.artifacttypes.myFileType:
    derived_from: tosca.artifacts.File
```

2004 **3.9.3.11 data_types**

2005 This optional keyname provides a section to define new data types in TOSCA.

2006 **3.9.3.11.1 Keyname**

```
data_types
```

2007 **3.9.3.11.2 Grammar**

```
data_types:
  <tosca_datatype_def_1>
  ...
  <tosca_datatype_def_n>
```

2008 **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
```

2009 3.9.3.12 capability_types

2010 This optional keyname lists the Capability Types that provide the reusable type definitions that can be used to
2011 describe features Node Templates or Node Types can declare they support.

2012 3.9.3.12.1 Keyname

```
capability_types
```

2013 3.9.3.12.2 Grammar

```
capability_types:
  <capability_type_defn_1>
  ...
  <capability_type_defn_n>
```

2014 3.9.3.12.3 Example

```
capability_types:
  mycompany.mytypes.myCustomEndpoint:
    derived_from: tosca.capabilities.Endpoint
    properties:
      # more details ...

  mycompany.mytypes.myCustomFeature:
    derived_from: tosca.capabilities.Feature
    properties:
      # more details ...
```

2015 3.9.3.13 interface_types

2016 This optional keyname lists the Interface Types that provide the reusable type definitions that can be used to
2017 describe operations for on TOSCA entities such as Relationship Types and Node Types.

2018 3.9.3.13.1 Keyname

```
interface_types
```

2019 3.9.3.13.2 Grammar

```
interface_types:
  <interface_type_defn_1>
  ...
  <interface type_defn_n>
```

2020 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.
```

2021 3.9.3.14 relationship_types

2022 This optional keyname lists the Relationship Types that provide the reusable type definitions that can be used to
2023 describe dependent relationships between Node Templates or Node Types.

2024 3.9.3.14.1 Keyname

```
relationship_types
```

2025 3.9.3.14.2 Grammar

```
relationship_types:
  <relationship_type_defn_1>
  ...
  <relationship type_defn_n>
```

2026 3.9.3.14.3 Example

```
relationship_types:
  mycompany.mytypes.myCustomClientServerType:
    derived_from: tosca.relationships.HostedOn
    properties:
      # more details ...

  mycompany.mytypes.myCustomConnectionType:
    derived_from: tosca.relationships.ConnectsTo
    properties:
```

```
# more details ...
```

2027 3.9.3.15 node_types

2028 This optional keyname lists the Node Types that provide the reusable type definitions for software components
2029 that Node Templates can be based upon.

2030 3.9.3.15.1 Keyname

```
node_types
```

2031 3.9.3.15.2 Grammar

```
node_types:  
  <node_type_defn_1>  
  ...  
  <node_type_defn_n>
```

2032 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
```

2033 3.9.3.15.4 Notes

- 2034
- The node types listed as part of the **node_types** block can be mapped to the list of **NodeType** definitions as described by the TOSCA v1.0 specification.
- 2035

2036 3.9.3.16 group_types

2037 This optional keyname lists the Group Types that are defined by this Service Template.

2038 3.9.3.16.1 Keyname

```
group_types
```

2039 3.9.3.16.2 Grammar

```
group_types:  
  <group_type_defn_1>
```



```
...  
<group type_defn_n>
```

2040 3.9.3.16.3 Example

```
group_types:  
  mycompany.mytypes.myScalingGroup:  
    derived_from: tosa.groups.Root
```

2041 3.9.3.17 policy_types

2042 This optional keyname lists the Policy Types that are defined by this Service Template.

2043 3.9.3.17.1 Keyname

```
policy_types
```

2044 3.9.3.17.2 Grammar

```
policy_types:  
  <policy_type_defn_1>  
  ...  
  <policy type_defn_n>
```

2045 3.9.3.17.3 Example

```
policy_types:  
  mycompany.mytypes.myScalingPolicy:  
    derived_from: tosa.policies.Scaling
```

2046
2047

2048
2049
2050
2051
2052
2053

2054

2055

2056
2057
2058
2059
2060
2061
2062
2063

4 TOSCA functions

This section is **normative** and includes functions that are supported for use within a TOSCA Service Template.

4.1 Reserved Function Keywords

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.

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	<p>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).</p> <p>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.</p>

4.2 Environment Variable Conventions

4.2.1 Reserved Environment Variable Names and Usage

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

The following keywords are reserved environment variable names in any TOSCA supported execution environment:

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).

Keyword	Valid Contexts	Description
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).

2064

2065 For scripts (or implementation artifacts in general) that run in the context of relationship operations, select
2066 properties and attributes of both the relationship itself as well as select properties and attributes of the source and
2067 target node(s) of the relationship can be provided to the environment by declaring respective operation inputs.

2068

2069 Declared inputs from mapped properties or attributes of the source or target node (selected via the **SOURCE** or
2070 **TARGET** keyword) will be provided to the environment as variables having the exact same name as the inputs. In
2071 addition, the same values will be provided for the complete set of source or target nodes, however prefixed with
2072 the ID of the respective nodes. By means of the **SOURCES** or **TARGETS** variables holding the complete set of source
2073 or target node IDs, scripts will be able to iterate over corresponding inputs for each provided ID prefix.

2074

2075 The following example snippet shows an imaginary relationship definition from a load-balancer node to worker
2076 nodes. A script is defined for the **add_target** operation of the Configure interface of the relationship, and the
2077 **ip_address** attribute of the target is specified as input to the script:

2078

```
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
```

2079 The **add_target** operation will be invoked, whenever a new target member is being added to the load-balancer.
2080 With the above inputs declaration, a **member_ip** environment variable that will hold the IP address of the target

2081 being added will be provided to the **configure_members.py** script. In addition, the IP addresses of all current
2082 load-balancer members will be provided as environment variables with a naming scheme of **<target node**
2083 **ID>_member_ip**. This will allow, for example, scripts that always just write the complete list of load-balancer
2084 members into a configuration file to do so instead of updating existing list, which might be more complicated.

2085 Assuming that the TOSCA application instance includes five load-balancer members, **node1** through **node5**,
2086 where **node5** is the current target being added, the following environment variables (plus potentially more
2087 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
```

2088 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 ...
```

2089 4.2.2 Prefixed vs. Unprefixed TARGET names

2090 The list target node types assigned to the TARGETS key in an execution environment would have names prefixed
2091 by unique IDs that distinguish different instances of a node in a running model. Future drafts of this specification
2092 will show examples of how these names/IDs will be expressed.

2093 4.2.2.1 Notes

- 2094 • Target of interest is always un-prefixed. Prefix is the target opaque ID. The IDs can be used to find the
2095 environment var. for the corresponding target. Need an example here.
- 2096 • If you have one node that contains multiple targets this would also be used (add or remove target
2097 operations would also use this you would get set of all current targets).

2098 4.3 Intrinsic functions

2099 These functions are supported within the TOSCA template for manipulation of template data.

2100 **4.3.1 concat**

2101 The **concat** function is used to concatenate two or more string values within a TOSCA service template.

2102 **4.3.1.1 Grammar**

```
concat: [<string_value_expressions_*> ]
```

2103 **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.

2104 **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 ] ] }
```

2105 **4.3.2 token**

2106 The **token** function is used within a TOSCA service template on a string to parse out (tokenize) substrings
2107 separated by one or more token characters within a larger string.

2108 **4.3.2.1 Grammar**

```
token: [ <string_with_tokens>, <string_of_token_chars>, <substring_index> ]
```

2109 **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.

2110 **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 ] }
```

4.4 Property functions

These functions are used within a service template to obtain property values from property definitions declared elsewhere in the same service template. These property definitions can appear either directly in the service template itself (e.g., in the **inputs** section) or on entities (e.g., node or relationship templates) that have been modeled within the template.

Note that the **get_input** and **get_property** functions may only retrieve the static values of property definitions of a TOSCA application as defined in the TOSCA Service Template. The **get_attribute** function should be used to retrieve values for attribute definitions (or property definitions reflected as attribute definitions) from the runtime instance model of the TOSCA application (as realized by the TOSCA orchestrator).

4.4.1 get_input

The **get_input** function is used to retrieve the values of properties declared within the **inputs** section of a TOSCA Service Template.

4.4.1.1 Grammar

```
get_input: <input_property_name>
```

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.

4.4.1.3 Examples

```
inputs:
  cpus:
    type: integer

node_templates:
  my_server:
    type: tosa.nodes.Compute
    capabilities:
      host:
        properties:
          num_cpus: { get_input: cpus }
```

4.4.2 get_property

The **get_property** function is used to retrieve property values between modelable entities defined in the same service template.

2130 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> ]
```

2131 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.

2132 4.4.2.3 Examples

2133 The following example shows how to use the get_property function with an actual Node Template name:

```
node_templates:  
  
  mysql_database:  
    type: toska.nodes.Database  
    properties:  
      name: sql_database1  
  
  wordpress:  
    type: toska.nodes.WebApplication.WordPress  
    ...  
    interfaces:  
      Standard:  
        configure:  
          inputs:  
            wp_db_name: { get_property: [ mysql_database, name ] }
```

2134 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 ] }

```

2135 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 ] }

```

2136 4.5 Attribute functions

2137 These functions (attribute functions) are used within an instance model to obtain attribute values from instances of
 2138 nodes and relationships that have been created from an application model described in a service template. The
 2139 instances of nodes or relationships can be referenced by their name as assigned in the service template or
 2140 relative to the context where they are being invoked.

2141 4.5.1 `get_attribute`

2142 The `get_attribute` function is used to retrieve the values of named attributes declared by the referenced node
 2143 or relationship template name.

2144

2145 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>, ]
```

2146 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.

2147 4.5.1.3 Examples:

2148 The attribute functions are used in the same way as the equivalent Property functions described above. Please
2149 see their examples and replace “get_property” with “get_attribute” function name.

2150 4.5.1.4 Notes

2151 These functions are used to obtain attributes from instances of node or relationship templates by the names they
2152 were given within the service template that described the application model (pattern).

- 2153 • These functions only work when the orchestrator can resolve to a single node or relationship instance for
2154 the named node or relationship. This essentially means this is acknowledged to work only when the node
2155 or relationship template being referenced from the service template has a cardinality of 1 (i.e., there can
2156 only be one instance of it running).

2157 4.6 Operation functions

2158 These functions are used within an instance model to obtain values from interface operations. These can be used
2159 in order to set an attribute of a node instance at runtime or to pass values from one operation to another.

2160 4.6.1 get_operation_output

2161 The **get_operation_output** function is used to retrieve the values of variables exposed / exported from an
2162 interface operation.

2163 4.6.1.1 Grammar

```
get_operation_output: <modelable_entity_name>, <interface_name>, <operation_name>,
<output_variable_name>
```

2164 **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.

2165 **4.6.1.3 Notes**

- 2166
- 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.
- 2167
- 2168

2169 **4.7 Navigation functions**

- 2170
- This version of the TOSCA Simple Profile does not define any model navigation functions.

2171 **4.7.1 get_nodes_of_type**

2172 The **get_nodes_of_type** function can be used to retrieve a list of all known instances of nodes of the declared Node Type.

2173

2174 **4.7.1.1 Grammar**

```
get_nodes_of_type: <node_type_name>
```

2175 **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.

2176 **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.

2177 **4.8 Artifact functions**

2178 **4.8.1 get_artifact**

2179 The **get_artifact** function is used to retrieve artifact location between modelable entities defined in the same service template.

2180

2181 **4.8.1.1 Grammar**

```
get_artifact: [ <modelable_entity_name>, <artifact_name>, <location>, <remove> ]
```

2182 **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	<p>Location value must be either a valid path e.g. '/etc/var/my_file' or 'LOCAL_FILE'.</p> <p>If the value is LOCAL_FILE the orchestrator is responsible for providing a path as the result of the get_artifact call where the artifact file can be accessed. The orchestrator will also remove the artifact from this location at the end of the operation.</p> <p>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 get_artifact function and leave the file here after the execution of the operation.</p>
remove	no	boolean	<p>Boolean flag to override the orchestrator default behavior so it will remove or not the artifact at the end of the operation execution.</p> <p>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.</p> <p>If true the artifact will be removed by the orchestrator at the end of the operation execution, if false it will not be removed.</p>

2183 **4.8.1.3 Examples**

2184 The following example shows how to use the **get_artifact** function with an actual Node Template name:

2185 **4.8.1.3.1 Example: Retrieving artifact without specified location:**

```
node_templates:

  wordpress:
    type: toska.nodes.WebApplication.WordPress
    ...
    interfaces:
      Standard:
        configure:
          create:
            implementation: wordpress_install.sh
            inputs
              wp_zip: { get_artifact: [ SELF, zip ] }
```

```
artifacts:
  zip: /data/wordpress.zip
```

In such implementation the TOSCA orchestrator may provide the wordpress.zip archive as a local URL (example: <file:///home/user/wordpress.zip>) or a remote one (example: <http://cloudrepo:80/files/wordpress.zip>) (some orchestrator may indeed provide some global artifact repository management features)

4.8.1.3.2 Example: Retrieving artifact as a local path :

The following example explains how to force the orchestrator to copy the file locally before calling the operation's implementation script:

```
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
```

In such implementation the tosca orchestrator must provide the wordpress.zip archive as a local path (example: </tmp/wordpress.zip>) and **will remove it** after the operation is completed.

4.8.1.3.3 Example: Retrieving artifact in a specified location:

The following example explains how to force the orchestrator to copy the file locally to a specific location before calling the operation's implementation script :

```
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
```

2199 In such implementation the tosca orchestrator must provide the wordpress.zip archive as a local path (example:
2200 C:/wpdata/wp.zip) and **will let it** after the operation is completed.

2201 4.9 Context-based Entity names (global)

2202 Future versions of this specification will address methods to access entity names based upon the context in which
2203 they are declared or defined.

2204 4.9.1.1 Goals

- 2205 • Using the full paths of modelable entity names to qualify context with the future goal of a more robust
2206 get_attribute function: e.g., get_attribute(<context-based-entity-name>, <attribute name>)

2207

5 TOSCA normative type definitions

2208 This section is **normative** and contains normative type definitions.

2209

2210 The declarative approach is heavily dependent of the definition of basic types that a declarative container
2211 must understand. The definition of these types must be very clear such that the operational semantics can
2212 be precisely followed by a declarative container to achieve the effects intended by the modeler of a
2213 topology in an interoperable manner.

2214

5.1 Assumptions

- 2215 • Assumes alignment with/dependence on XML normative types proposal for TOSCA v1.1
- 2216 • Assumes that the normative types will be versioned and the TOSCA TC will preserve backwards
2217 compatibility.
- 2218 • Assumes that security and access control will be addressed in future revisions or versions of this
2219 specification.

2220

5.2 Data Types

2221

5.2.1 `tosca.datatypes.Root`

2222 This is the default (root) TOSCA Root Type definition that all complex TOSCA Data Types derive from.

2223

5.2.1.1 Definition

2224 The TOSCA Credential type is defined as follows:

```
tosca.datatypes.Root:  
  description: The TOSCA root Data Type all other TOSCA base Data Types derive from
```

2225

5.2.2 `tosca.datatypes.Credential`

2226 The Credential type is a complex TOSCA data Type used when describing authorization credentials used to
2227 access network accessible resources.

Shorthand Name	Credential
Type Qualified Name	tosca:Credential
Type URI	tosca.datatypes.Credential

2228

5.2.2.1 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.

Name	Required	Type	Constraints	Description
userh	no	string	None	The optional user (name or ID) used for non-token based credentials.

5.2.2.2 Definition

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
```

5.2.2.3 Additional requirements

- TOSCA Orchestrators SHALL interpret and validate the value of the **token** property based upon the value of the **token_type** property.

5.2.2.4 Notes

- Specific token types and encoding them using network protocols are not defined or covered in this specification.
- The use of transparent user names (IDs) or passwords are not considered best practice.

5.2.2.5 Examples

5.2.2.5.1 Provide a simple user name and password without a protocol or standardized token format

```
<some_tosca_entity>:
  properties:
    my_credential:
      type: Credential
```

```
properties:
  user: myusername
  token: mypassword
```

2241 5.2.2.5.2 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
```

2242 5.2.2.5.3 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
```

2243 5.2.2.5.4 OAuth bearer token credential

```
<some_tosca_entity>:
  properties:
    my_credential: # type: Credential
      protocol: oauth2
      token_type: bearer
      # token encoded in Base64
      token: 8ao9nE2DEjr1zCsicWmpBC
```

2244 5.2.3 **tosca.datatypes.network.NetworkInfo**

2245 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

2246 **5.2.3.1 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.

2247 **5.2.3.2 Definition**

2248 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
```

2249 **5.2.3.3 Examples**

2250 Example usage of the NetworkInfo data type:

```
private_network:
  network_name: private
  network_id: 3e54214f-5c09-1bc9-9999-44100326da1b
  addresses: [ 10.111.128.10 ]
```

2251 **5.2.3.4 Additional Requirements**

- 2252
- 2253
- 2254
- It is expected that TOSCA orchestrators **MUST** be able to map the **network_name** from the 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.

2255 **5.2.4 toska.datatypes.network.PortInfo**

2256 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

2257 **5.2.4.1 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.

2258 **5.2.4.2 Definition**

2259 The TOSCA Port type is defined as follows:

```
tosca.datatypes.network.PortInfo:
  derived_from: toska.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
```

2260 **5.2.4.3 Examples**

2261 Example usage of the PortInfo data type:

```
ethernet_port:
  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 ]
```

2262 **5.2.4.4 Additional Requirements**

- 2263
- It is expected that TOSCA orchestrators MUST be able to map the **port_name** from the TOSCA model to underlying network model of the provider.
- 2264
- The properties (or attributes) of PortInfo may or may not be required depending on usage context.
- 2265

2266 **5.2.5 tosca.datatypes.network.PortDef**

2267 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

2268 **5.2.5.1 Definition**

2269 The TOSCA PortDef type is defined as follows:

```
tosca.datatypes.network.PortDef:
  derived_from: integer
  constraints:
    - in_range: [ 1, 65535 ]
```

2270 **5.2.5.2 Examples**

2271 Example use of a PortDef property type:

```
listen_port:
  type: PortDef
  default: 9000
  constraints:
    - in_range: [ 9000, 9090 ]
```

2272 **5.2.6 tosca.datatypes.network.PortSpec**

2273 The PortSpec type is a complex TOSCA data Type used when describing port specifications for a network connection.

2274

Shorthand Name	PortSpec
Type Qualified Name	tosca:PortSpec
Type URI	tosca.datatypes.network.PortSpec

2275 **5.2.6.1 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.
source_range	no	range	in_range: [1, 65536]	The optional range for source port.
target	no	PortDef	See PortDef	The optional target port.

Name	Required	Type	Constraints	Description
target_range	no	range	in_range: [1, 65536]	The optional range for target port.

5.2.6.2 Definition

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
    target_range:
      type: range
      constraints:
        - in_range: [ 1, 65535 ]
    source:
      type: PortDef
    source_range:
      type: range
      constraints:
        - in_range: [ 1, 65535 ]
```

5.2.6.3 Additional requirements

- A valid PortSpec must have at least one of the following properties: **target**, **target_range**, **source** or **source_range**.

5.2.6.4 Examples

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 ]
```

5.3 Artifact Types

TOSCA Artifacts Types represent the types of packages and files used by the orchestrator when deploying TOSCA Node or Relationship Types or invoking their interfaces. Currently, artifacts are logically divided into three categories:

- **Deployment Types:** includes those artifacts that are used during deployment (e.g., referenced on create and install operations) and include packaging files such as RPMs, ZIPs, or TAR files.
- **Implementation Types:** includes those artifacts that represent imperative logic and are used to implement TOSCA Interface operations. These typically include scripting languages such as Bash (.sh), Chef and Puppet.
- **Runtime Types:** includes those artifacts that are used during runtime by a service or component of the application. This could include a library or language runtime that is needed by an application such as a PHP or Java library.

Note: Additional TOSCA Artifact Types will be developed in future drafts of this specification.

5.3.1 `tosca.artifacts.Root`

This is the default (root) TOSCA [Artifact Type](#) definition that all other TOSCA base Artifact Types derive from.

5.3.1.1 Definition

```
tosca.artifacts.Root:
  description: The TOSCA Artifact Type all other TOSCA Artifact Types derive from
```

5.3.2 `tosca.artifacts.File`

This artifact type is used when an artifact definition needs to have its associated file simply treated as a file and no special handling/handlers are invoked (i.e., it is not treated as either an implementation or deployment artifact type).

Shorthand Name	File
Type Qualified Name	tosca:File
Type URI	tosca.artifacts.File

5.3.2.1 Definition

```
tosca.artifacts.File:
  derived_from: tosca.artifacts.Root
```

5.3.3 Deployment Types

5.3.3.1 `tosca.artifacts.Deployment`

This artifact type represents the parent type for all deployment artifacts in TOSCA. This class of artifacts typically represents a binary packaging of an application or service that is used to install/create or deploy it as part of a node's lifecycle.

2311 **5.3.3.1.1 Definition**

```
tosca.artifacts.Deployment:
  derived_from: tosca.artifacts.Root
  description: TOSCA base type for deployment artifacts
```

2312 **5.3.3.2 Additional Requirements**

- 2313 • TOSCA Orchestrators MAY throw an error if it encounters a non-normative deployment artifact type that it
2314 is not able to process.

2315 **5.3.3.3 [tosca.artifacts.Deployment.Image](#)**

2316 This artifact type represents a parent type for any “image” which is an opaque packaging of a TOSCA Node’s
2317 deployment (whether real or virtual) whose contents are typically already installed and pre-configured (i.e.,
2318 “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

2319 **5.3.3.3.1 Definition**

```
tosca.artifacts.Deployment.Image:
  derived_from: tosca.artifacts.Deployment
```

2320 **5.3.3.4 [tosca.artifacts.Deployment.Image.VM](#)**

2321 This artifact represents the parent type for all Virtual Machine (VM) image and container formatted deployment
2322 artifacts. These images contain a stateful capture of a machine (e.g., server) including operating system and
2323 installed software along with any configurations and can be run on another machine using a hypervisor which
2324 virtualizes typical server (i.e., hardware) resources.

2325 **5.3.3.4.1 Definition**

```
tosca.artifacts.Deployment.Image.VM:
  derived_from: tosca.artifacts.Deployment.Image
  description: Virtual Machine (VM) Image
```

2326 **5.3.3.4.2 Notes**

- 2327 • Future drafts of this specification may include popular standard VM disk image (e.g., ISO, VMI, VMDX,
2328 QCOW2, etc.) and container (e.g., OVF, bare, etc.) formats. These would include consideration of disk
2329 formats such as:

2330 **5.3.4 Implementation Types**

2331 **5.3.4.1 [tosca.artifacts.Implementation](#)**

2332 This artifact type represents the parent type for all implementation artifacts in TOSCA. These artifacts are used to
2333 implement operations of TOSCA interfaces either directly (e.g., scripts) or indirectly (e.g., config. files).

2334 **5.3.4.1.1 Definition**

```
tosca.artifacts.Implementation:
  derived_from: tosca.artifacts.Root
  description: TOSCA base type for implementation artifacts
```

2335 **5.3.4.2 Additional Requirements**

- 2336 • TOSCA Orchestrators **MAY** throw an error if it encounters a non-normative implementation artifact type
2337 that it is not able to process.

2338 **5.3.4.3 [tosca.artifacts.Implementation.Bash](#)**

2339 This artifact type represents a Bash script type that contains Bash commands that can be executed on the Unix
2340 Bash shell.

Shorthand Name	Bash
Type Qualified Name	tosca:Bash
Type URI	tosca.artifacts.Implementation.Bash

2341 **5.3.4.3.1 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 ]
```

2342 **5.3.4.4 [tosca.artifacts.Implementation.Python](#)**

2343 This artifact type represents a Python file that contains Python language constructs that can be executed within a
2344 Python interpreter.

Shorthand Name	Python
Type Qualified Name	tosca:Python
Type URI	tosca.artifacts.Implementation.Python

2345 **5.3.4.4.1 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 ]
```

2346

5.4 Capabilities Types

2347

5.4.1 `tosca.capabilities.Root`

2348 This is the default (root) TOSCA Capability Type definition that all other TOSCA Capability Types derive from.

2349

5.4.1.1 Definition

```
tosca.capabilities.Root:
  description: The TOSCA root Capability Type all other TOSCA base Capability Types
  derive from
```

2350

5.4.2 `tosca.capabilities.Node`

2351 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

2352

5.4.2.1 Definition

```
tosca.capabilities.Node:
  derived_from: tosca.capabilities.Root
```

2353

5.4.3 `tosca.capabilities.Container`

2354 The Container capability, when included on a Node Type or Template definition, indicates that the node can act
2355 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

2356

5.4.3.1 Properties

Name	Required	Type	Constraints	Description
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).

2357


```
tosca.capabilities.Container:
  derived_from: toska.capabilities.Root
  properties:
    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
```

2359 5.4.4 toska.capabilities.Endpoint

2360 This is the default TOSCA type that should be used or extended to define a network endpoint capability. This
2361 includes the information to express a basic endpoint with a single port or a complex endpoint with multiple ports.
2362 By default the Endpoint is assumed to represent an address on a private network unless otherwise specified.

Shorthand Name	Endpoint
Type Qualified Name	tosca:Endpoint
Type URI	tosca.capabilities.Endpoint

2363 5.4.4.1 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.

Name	Required	Type	Constraints	Description
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: <ul style="list-style-type: none"> 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)

2364 5.4.4.2 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.

2365 5.4.4.3 Definition

```

tosca.capabilities.Endpoint:
  derived_from: tosca.capabilities.Root
  properties:
    protocol:
      type: string
      default: tcp
    port:
      type: PortDef
      required: false
    secure:
      type: boolean
      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
  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

```

5.4.4.4 Additional requirements

- Although both the port and ports properties are not required, one of port or ports must be provided in a valid [Endpoint](#).

5.4.5 **tosca.capabilities.Endpoint.Public**

This capability represents a public endpoint which is accessible to the general internet (and its public IP address ranges).

This public endpoint capability also can be used to create a floating (IP) address that the underlying network assigns from a pool allocated from the application's underlying public network. This floating address is managed by the underlying network such that can be routed an application's private address and remains reliable to internet clients.

Shorthand Name	Endpoint.Public
Type Qualified Name	tosca:Endpoint.Public
Type URI	tosca.capabilities.Endpoint.Public

5.4.5.1 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: 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

```

5.4.5.2 Additional requirements

- If the **network_name** is set to the reserved value **PRIVATE** or if the value is set to the name of network (or subnetwork) that is not public (i.e., has non-public IP address ranges assigned to it) then TOSCA Orchestrators **SHALL** treat this as an error.
- If a **dns_name** is set, TOSCA Orchestrators SHALL attempt to register the name in the (local) DNS registry for the Cloud provider.

5.4.6 **tosca.capabilities.Endpoint.Admin**

This is the default TOSCA type that should be used or extended to define a specialized administrator endpoint capability.

Shorthand Name	Endpoint.Admin
Type Qualified Name	tosca:Endpoint.Admin
Type URI	tosca.capabilities.Endpoint.Admin

5.4.6.1 Properties

Name	Required	Type	Constraints	Description
None	N/A	N/A	N/A	N/A

5.4.6.2 Definition

```

tosca.capabilities.Endpoint.Admin:
  derived_from: tosca.capabilities.Endpoint
  # Change Endpoint secure indicator to true from its default of false
  properties:
    secure: true

```

5.4.6.3 Additional requirements

- TOSCA Orchestrator implementations of Endpoint.Admin (and connections to it) **SHALL** assure that network-level security is enforced if possible.

5.4.7 **tosca.capabilities.Endpoint.Database**

This is the default TOSCA type that should be used or extended to define a specialized database endpoint capability.

Shorthand Name	Endpoint.Database
Type Qualified Name	tosca:Endpoint.Database
Type URI	tosca.capabilities.Endpoint.Database

5.4.7.1 Properties

Name	Required	Type	Constraints	Description
None	N/A	N/A	N/A	N/A

5.4.7.2 Definition

```
tosca.capabilities.Endpoint.Database:
  derived_from: tosca.capabilities.Endpoint
```

5.4.8 tosca.capabilities.Attachment

This is the default TOSCA type that should be used or extended to define an attachment capability of a (logical) infrastructure device node (e.g., [BlockStorage](#) node).

Shorthand Name	Attachment
Type Qualified Name	tosca:Attachment
Type URI	tosca.capabilities.Attachment

5.4.8.1 Properties

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

5.4.8.2 Definition

```
tosca.capabilities.Attachment:
  derived_from: tosca.capabilities.Root
```

5.4.9 tosca.capabilities.OperatingSystem

This is the default TOSCA type that should be used to express an Operating System capability for a node.

Shorthand Name	OperatingSystem
Type Qualified Name	tosca:OperatingSystem
Type URI	tosca.capabilities.OperatingSystem

5.4.9.1 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.

Name	Required	Type	Constraints	Description
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.

5.4.9.2 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

```

5.4.9.3 Additional Requirements

- Please note that the string values for the properties **architecture**, **type** and **distribution** SHALL be normalized to lowercase by processors of the service template for matching purposes. For example, if a “**type**” value is set to either “Linux”, “LINUX” or “linux” in a service template, the processor would normalize all three values to “linux” for matching purposes.

5.4.9.4 Notes

- None

5.4.10 **tosca.capabilities.Scalable**

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

5.4.10.1 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.

5.4.10.2 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

```

5.4.10.3 Notes

- The actual number of instances for a node may be governed by a separate scaling policy which conceptually would be associated to either a scaling-capable node or a group of nodes in which it is defined to be a part of. This is a planned future feature of the TOSCA Simple Profile and not currently described.

5.4.11 tosca.capabilities.network.Bindable

A node type that includes the Bindable capability indicates that it can be bound to a logical network association via a network port.

Shorthand Name	network.Bindable
Type Qualified Name	tosca:network.Bindable
Type URI	tosca.capabilities.network.Bindable

5.4.11.1 Properties

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

5.4.11.2 Definition

```
tosca.capabilities.network.Bindable:
  derived_from: tosca.capabilities.Node
```

5.5 Requirement Types

There are no normative Requirement Types currently defined in this working draft. Typically, Requirements are described against a known Capability Type

5.6 Relationship Types

5.6.1 [tosca.relationships.Root](#)

This is the default (root) TOSCA Relationship Type definition that all other TOSCA Relationship Types derive from.

5.6.1.1 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.

5.6.1.2 Definition

```
tosca.relationships.Root:
  description: The TOSCA root Relationship Type all other TOSCA base Relationship
Types derive from
  attributes:
    toska_id:
      type: string
    toska_name:
      type: string
```



```
interfaces:
  Configure:
    type: tosca.interfaces.relationship.Configure
```

2435 5.6.2 [tosca.relationships.DependsOn](#)

2436 This type represents a general dependency relationship between two nodes.

Shorthand Name	DependsOn
Type Qualified Name	tosca:DependsOn
Type URI	tosca.relationships.DependsOn

2437 5.6.2.1 Definition

```
tosca.relationships.DependsOn:
  derived_from: tosca.relationships.Root
  valid_target_types: [ tosca.capabilities.Node ]
```

2438 5.6.3 [tosca.relationships.HostedOn](#)

2439 This type represents a hosting relationship between two nodes.

Shorthand Name	HostedOn
Type Qualified Name	tosca:HostedOn
Type URI	tosca.relationships.HostedOn

2440 5.6.3.1 Definition

```
tosca.relationships.HostedOn:
  derived_from: tosca.relationships.Root
  valid_target_types: [ tosca.capabilities.Container ]
```

2441 5.6.4 [tosca.relationships.ConnectsTo](#)

2442 This type represents a network connection relationship between two nodes.

Shorthand Name	ConnectsTo
Type Qualified Name	tosca:ConnectsTo
Type URI	tosca.relationships.ConnectsTo

2443 5.6.4.1 Definition

```
tosca.relationships.ConnectsTo:
  derived_from: tosca.relationships.Root
  valid_target_types: [ tosca.capabilities.Endpoint ]
  properties:
    credential:
```

```
type: tosca.datatypes.Credential
required: false
```

2444 5.6.4.2 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.

2445 5.6.5 [tosca.relationships.AttachesTo](#)

2446 This type represents an attachment relationship between two nodes. For example, an AttachesTo relationship
2447 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

2448 5.6.5.1 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'

2449 5.6.5.2 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.

2450 5.6.5.3 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

5.6.6 `tosca.relationships.RoutesTo`

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

5.6.6.1 Definition

```
tosca.relationships.RoutesTo:
  derived_from: tosca.relationships.ConnectsTo
  valid_target_types: [ tosca.capabilities.Endpoint ]
```

5.7 Interface Types

Interfaces are reusable entities that define a set of operations that that can be included as part of a Node type or Relationship Type definition. Each named operations may have code or scripts associated with them that orchestrators can execute for when transitioning an application to a given state.

5.7.1 Additional Requirements

- Designers of Node or Relationship types are not required to actually provide/associate code or scripts with every operation for a given interface it supports. In these cases, orchestrators SHALL consider that a “No Operation” or “no-op”.
- The default behavior when providing scripts for an operation in a sub-type (sub-class) or a 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.

5.7.2 Best Practices

- When TOSCA Orchestrators substitute an implementation for an abstract node in a deployed service template it SHOULD be able to present a confirmation to the submitter to confirm the implementation chosen would be acceptable.

5.7.3 `tosca.interfaces.node.lifecycle.Standard`

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

5.7.3.1 Definition

```
tosca.interfaces.node.lifecycle.Standard:
  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.

```

5.7.3.2 Create operation

The create operation is generally used to create the resource or service the node represents in the topology. TOSCA orchestrators expect node templates to provide either a deployment artifact or an implementation artifact of a defined artifact type that it is able to process. This specification defines normative deployment and implementation artifact types all TOSCA Orchestrators are expected to be able to process to support application portability.

5.7.3.3 Using TOSCA artifacts with interface operations

```

# Bash script whose location is described in the TOSCA CSAR file
# 1) LOCAL SCOPE: symbolic artifact name <or> 2) GLOBAL: actual file name from CSAR
#

```

5.7.3.4 TOSCA Orchestrator processing of Deployment artifacts

TOSCA Orchestrators, when encountering a deployment artifact on the create operation; will automatically attempt to deploy the artifact based upon its artifact type. This means that no implementation artifacts (e.g., scripts) are needed on the create operation to provide commands that deploy or install the software.

For example, if a TOSCA Orchestrator is processing an application with a node of type SoftwareComponent and finds that the node's template has a create operation that provides a filename (or references to an artifact which describes a file) of a known TOSCA deployment artifact type such as an Open Virtualization Format (OVF) image it will automatically deploy that image into the SoftwareComponent's host Compute node.

5.7.3.5 Operation sequencing and node state


The following diagrams show how TOSCA orchestrators sequence the operations of the Standard lifecycle in normal node startup and shutdown procedures.

The following key should be used to interpret the diagrams:



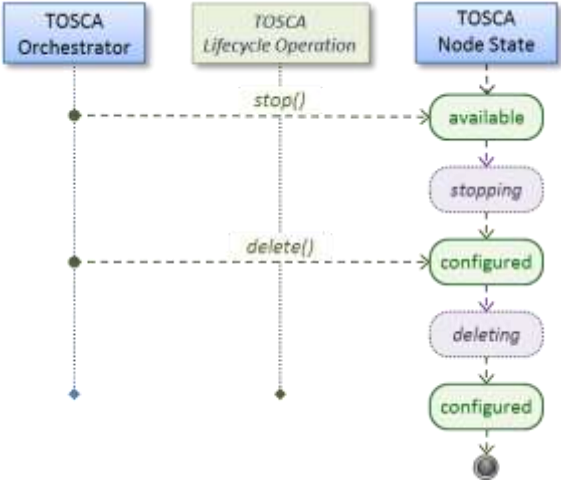
5.7.3.5.1 Normal node startup sequence diagram

2497 The following diagram shows how the TOSCA orchestrator would invoke operations on the Standard lifecycle to
2498 startup a node.



2499 **5.7.3.5.2 Normal node shutdown sequence diagram**

2500 The following diagram shows how the TOSCA orchestrator would invoke operations on the Standard lifecycle to
2501 shut down a node.



2502

2503 **5.7.4 tosca.interfaces.relationship.Configure**

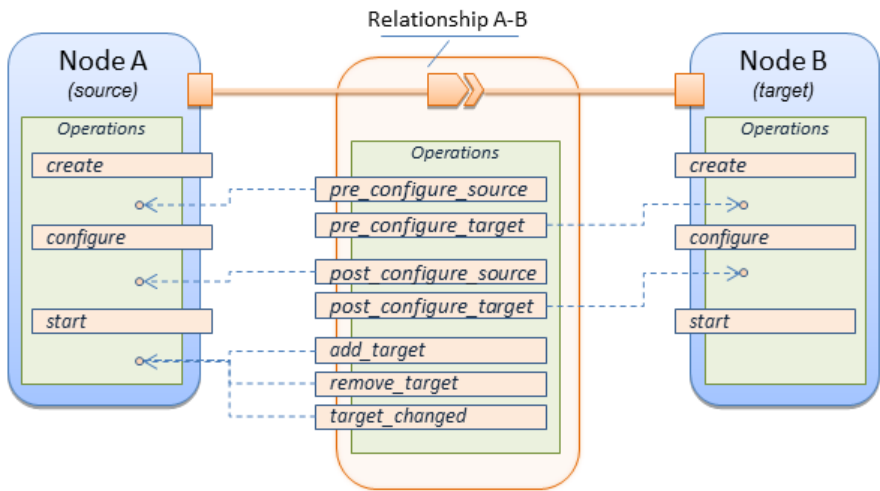
2504 The lifecycle interfaces define the essential, normative operations that each TOSCA Relationship Types may
2505 support.

Shorthand Name	Configure
Type Qualified Name	tosca:Configure
Type URI	tosca.interfaces.relationship.Configure

2506 **5.7.4.1 Definition**

```
tosca.interfaces.relationship.Configure:
  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.
```

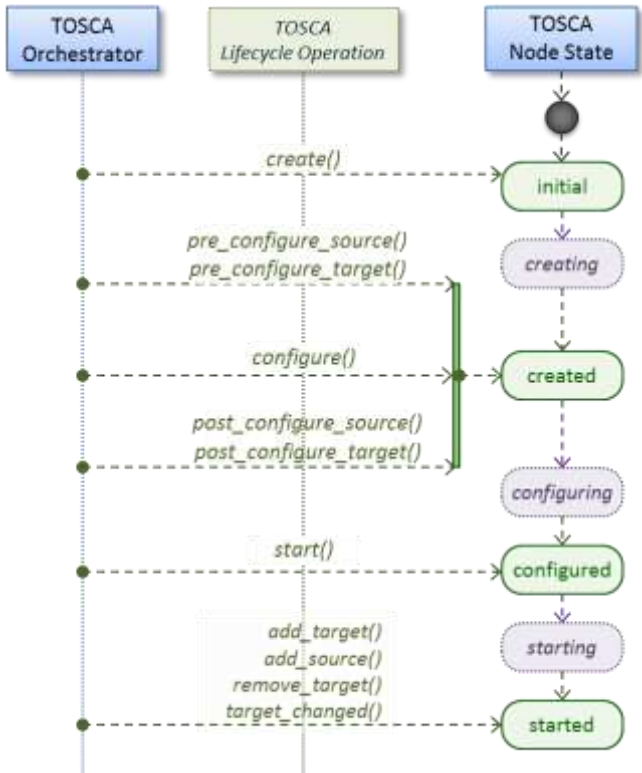
2507



2509 TOSCA relationships are directional connecting a source node to a target node. When TOSCA Orchestrator
2510 connects a source and target node together using a relationship that supports the Configure interface it will
2511 “interleave” the operations invocations of the Configure interface with those of the node’s own Standard lifecycle
2512 interface. This concept is illustrated below:

2513 5.7.4.3 Normal node start sequence with Configure relationship operations

2514 The following diagram shows how the TOSCA orchestrator would invoke Configure lifecycle operations in
2515 conjunction with Standard lifecycle operations during a typical startup sequence on a node.



2516 5.7.4.4 Node-Relationship configuration sequence

2517 Depending on which side (i.e., source or target) of a relationship a node is on, the orchestrator will:

2518 Invoke either the **pre_configure_source** or **pre_configure_target** operation as supplied by the
2519 relationship on the node.

2520 Invoke the node’s **configure** operation.

2521 Invoke either the **post_configure_source** or **post_configure_target** as supplied by the
2522 relationship on the node.

2523 Note that the **pre_configure_xxx** and **post_configure_xxx** are invoked only once per node instance.

2524 **5.7.4.4.1 Node-Relationship add, remove and changed sequence**

2525 Since a topology template contains nodes that can dynamically be added (and scaled), removed or changed as
2526 part of an application instance, the Configure lifecycle includes operations that are invoked on node instances that
2527 to notify and address these dynamic changes.

2528

2529 For example, a source node, of a relationship that uses the Configure lifecycle, will have the relationship
2530 operations **add_target**, or **remove_target** invoked on it whenever a target node instance is added or removed
2531 to the running application instance. In addition, whenever the node state of its target node changes, the
2532 **target_changed** operation is invoked on it to address this change. Conversely, the **add_source** and
2533 **remove_source** operations are invoked on the source node of the relationship.

2534 **5.7.4.5 Notes**

- 2535 • The target (provider) **MUST** be active and running (i.e., all its dependency stack **MUST** be fulfilled) prior to
2536 invoking **add_target**
 - 2537 • In other words, all Requirements **MUST** be satisfied before it advertises its capabilities (i.e., the
2538 attributes of the matched Capabilities are available).
 - 2539 • In other words, it cannot be “consumed” by any dependent node.
 - 2540 • Conversely, since the source (consumer) needs information (attributes) about any targets (and their
2541 attributes) being removed before it actually goes away.
- 2542 • The **remove_target** operation should only be executed if the target has had **add_target** executed. BUT
2543 in truth we’re first informed about a target in **pre_configure_source**, so if we execute that the source
2544 node should see **remove_target** called to cleanup.
- 2545 • **Error handling:** If any node operation of the topology fails processing should stop on that node template
2546 and the failing operation (script) should return an error (failure) code when possible.

2547 **5.8 Node Types**

2548 **5.8.1 toska.nodes.Root**

2549 The TOSCA **Root** Node Type is the default type that all other TOSCA base Node Types derive from. This allows
2550 for all TOSCA nodes to have a consistent set of features for modeling and management (e.g., consistent
2551 definitions for requirements, capabilities and lifecycle interfaces).

2552

Shorthand Name	Root
Type Qualified Name	tosca:Root
Type URI	tosca.nodes.Root

2553 **5.8.1.1 Properties**

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	The TOSCA Root Node type has no specified properties.

2554

5.8.1.2 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.

2555

5.8.1.3 Definition

```
tosca.nodes.Root:
  description: The TOSCA Node Type all other TOSCA base Node Types derive from
  attributes:
    toska_id:
      type: string
    toska_name:
      type: string
    state:
      type: string
  capabilities:
    feature:
      type: toska.capabilities.Node
  requirements:
    - dependency:
        capability: toska.capabilities.Node
        node: toska.nodes.Root
        relationship: toska.relationships.DependsOn
        occurrences: [ 0, UNBOUNDED ]
  interfaces:
    Standard:
      type: toska.interfaces.node.lifecycle.Standard
```

2556

5.8.1.4 Additional Requirements

- 2557
- 2558
- All Node Type definitions that wish to adhere to the TOSCA Simple Profile **SHOULD** extend from the TOSCA Root Node Type to be assured of compatibility and portability across implementations.

2559

5.8.2 toska.nodes.Compute

2560

2561

2562

The TOSCA **Compute** node represents one or more real or virtual processors of software applications or services along with other essential local resources. Collectively, the resources the compute node represents can logically be viewed as a (real or virtual) “server”.

Shorthand Name	Compute
Type Qualified Name	tosca:Compute
Type URI	tosca.nodes.Compute

2563 5.8.2.1 Properties

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

2564 5.8.2.2 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.

2565 5.8.2.3 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

```

5.8.2.4 Additional Requirements

- 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.

5.8.3 [tosca.nodes.SoftwareComponent](#)

The TOSCA **SoftwareComponent** node represents a generic software component that can be managed and run by a TOSCA **Compute** Node Type.

Shorthand Name	SoftwareComponent
Type Qualified Name	tosca:SoftwareComponent
Type URI	tosca.nodes.SoftwareComponent

5.8.3.1 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.

5.8.3.2 Attributes

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

5.8.3.3 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

```

5.8.3.4 Additional Requirements

- Nodes that can directly be managed and run by a TOSCA **Compute** Node Type **SHOULD** extend from this type.

5.8.4 tosca.nodes.WebServer

This TOSA **WebServer** Node Type represents an abstract software component or service that is capable 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

5.8.4.1 Properties

Name	Required	Type	Constraints	Description
None	N/A	N/A	N/A	N/A

5.8.4.2 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 ]

```

5.8.4.3 Notes and Additional Requirements

- This node **SHALL** export both a secure endpoint capability (i.e., **admin_endpoint**), typically for administration, as well as a regular endpoint (i.e., **data_endpoint**) for serving data.

2587 **5.8.5 tosca.nodes.WebApplication**

2588 The TOSCA **WebApplication** node represents a software application that can be managed and run by a TOSCA
2589 **WebServer** node. Specific types of web applications such as Java, etc. could be derived from this type.

Shorthand Name	WebApplication
Type Qualified Name	tosca: WebApplication
Type URI	tosca.nodes.WebApplication

2590 **5.8.5.1 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.

2591 **5.8.5.2 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
```

2592 **5.8.5.3 Additional Requirements**

- 2593
- None

2594 **5.8.6 tosca.nodes.DBMS**

2595 The TOSCA **DBMS** node represents a typical relational, SQL Database Management System software component
2596 or service.

2597 **5.8.6.1 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.

2598 **5.8.6.2 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 ]
```

2599 **5.8.6.3 Additional Requirements**

- 2600
- None

2601 **5.8.7 [tosca.nodes.Database](#)**

2602 The TOSCA **Database** node represents a logical database that can be managed and hosted by a TOSCA **DBMS**
2603 node.

Shorthand Name	Database
Type Qualified Name	tosca:Database
Type URI	tosca.nodes.Database

2604 **5.8.7.1 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.

```
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
```

2606 5.8.8 **tosca.nodes.ObjectStorage**

2607 The TOSCA **ObjectStorage** node represents storage that provides the ability to store data as objects (or BLOBs
2608 of data) without consideration for the underlying filesystem or devices.

Shorthand Name	ObjectStorage
Type Qualified Name	tosca:ObjectStorage
Type URI	tosca.nodes.ObjectStorage

2609 5.8.8.1 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).

```
tosca.nodes.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
```

2611 5.8.8.3 Notes:

- 2612
- Subclasses of the `tosca.nodes.ObjectStorage` node type may impose further constraints on
- 2613 properties. For example, a subclass may constrain the (minimum or maximum) length of the ‘**name**’
- 2614 property or include a regular expression to constrain allowed characters used in the ‘**name**’ property.

2615 5.8.9 [tosca.nodes.BlockStorage](#)

2616 The TOSCA **BlockStorage** node currently represents a server-local block storage device (i.e., not shared)

2617 offering evenly sized blocks of data from which raw storage volumes can be created.

2618 **Note:** In this draft of the TOSCA Simple Profile, distributed or Network Attached Storage (NAS) are not yet

2619 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.BlockStorage

2620 5.8.9.1 Properties

Name	Required	Type	Constraints	Description
size	yes *	scalar-unit.size	greater_or_equal : 1 MB	<p>The requested storage size (default unit is MB).</p> <p>* Note:</p> <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.

Name	Required	Type	Constraints	Description
volume_id	no	string	None	ID of an existing volume (that is in the accessible scope of the requesting application).
snapshot_id	no	string	None	Some identifier that represents an existing snapshot that should be used when creating the block storage (volume).

5.8.9.2 Attributes

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

5.8.9.3 Definition

```

tosca.nodes.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

```

5.8.9.4 Additional Requirements

- The **size** property is required when an existing volume (i.e., **volume_id**) is not available. However, if the property **volume_id** is provided, the **size** property is ignored.

5.8.9.5 Notes

- Resize of existing volumes is not considered at this time.
- 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).
- Currently, this version of the Simple Profile does not consider regions (or availability zones) when modeling storage.

2633 **5.8.10 `tosca.nodes.Container.Runtime`**

2634 The TOSCA **Container** Runtime node represents operating system-level virtualization technology used to run
2635 multiple application services on a single Compute host.

Shorthand Name	Container.Runtime
Type Qualified Name	tosca:Container.Runtime
Type URI	tosca.nodes.Container.Runtime

2636 **5.8.10.1 Definition**

```
tosca.nodes.Container.Runtime:
  derived_from: tosca.nodes.SoftwareComponent
  capabilities:
    host:
      type: tosca.capabilities.Container
  scalable:
    type: tosca.capabilities.Scalable
```

2637 **5.8.11 `tosca.nodes.Container.Application`**

2638 The TOSCA **Container** Application node represents an application that requires **Container**-level virtualization
2639 technology.

Shorthand Name	Container.Application
Type Qualified Name	tosca:Container.Application
Type URI	tosca.nodes.Container.Application

2640 **5.8.11.1 Definition**

```
tosca.nodes.Container.Application:
  derived_from: tosca.nodes.Root
  requirements:
    - host:
        capability: tosca.capabilities.Container
        node: tosca.nodes.Container
        relationship: tosca.relationships.HostedOn
```

2641 **5.8.12 `tosca.nodes.LoadBalancer`**

2642 The TOSCA **Load Balancer** node represents logical function that be used in conjunction with a Floating Address
2643 to distribute an application's traffic (load) across a number of instances of the application (e.g., for a clustered or
2644 scaled application).

Shorthand Name	LoadBalancer
Type Qualified Name	tosca:LoadBalancer
Type URI	tosca.nodes.LoadBalancer

5.8.12.1 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

```

5.8.12.2 Notes:

- 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 occurrences.

5.9 Group Types

TOSCA Group Types represent logical groupings of TOSCA nodes that have an implied membership relationship and may need to be orchestrated or managed together to achieve some result. Some use cases being developed by the TOSCA TC use groups to apply TOSCA policies for software placement and scaling while other use cases show groups can be used to describe cluster relationships.

Note: Additional normative TOSCA Group Types and use cases for them will be developed in future drafts of this specification.

5.9.1 tosca.groups.Root

This is the default (root) TOSCA [Group Type](#) definition that all other TOSCA base Group Types derive from.

5.9.1.1 Definition

```

tosca.groups.Root:
  description: The TOSCA Group Type all other TOSCA Group Types derive from

```

```
interfaces:
  Standard:
    type: tosca.interfaces.node.lifecycle.Standard
```

5.9.1.2 Notes:

- Group operations are not necessarily tied directly to member nodes that are part of a group.
- Future versions of this specification will create sub types of the **tosca.groups.Root** type that will describe how Group Type operations are to be orchestrated.

5.10 Policy Types

TOSCA Policy Types represent logical grouping of TOSCA nodes that have an implied relationship and need to be orchestrated or managed together to achieve some result. Some use cases being developed by the TOSCA TC use groups to apply TOSCA policies for software placement and scaling while other use cases show groups can be used to describe cluster relationships.

5.10.1 **tosca.policies.Root**

This is the default (root) TOSCA Policy Type definition that all other TOSCA base Policy Types derive from.

5.10.1.1 Definition

```
tosca.policies.Root:
  description: The TOSCA Policy Type all other TOSCA Policy Types derive from
```

5.10.1.2 Notes:

- None

5.10.2 **tosca.policies.Placement**

This is the default (root) TOSCA Policy Type definition that is used to govern placement of TOSCA nodes or groups of nodes.

5.10.2.1 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.
```

5.10.2.2 Notes:

- None

5.10.3 **tosca.policies.Scoring**

This is the default (root) TOSCA Policy Type definition that is used to govern scaling of TOSCA nodes or groups of nodes.

2683 **5.10.3.1 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.
```

2684 **5.10.3.2 Notes:**

- 2685
- None

2686 **5.10.4 tosca.policies.Update**

2687 This is the default (root) TOSCA Policy Type definition that is used to govern update of TOSCA nodes or groups
2688 of nodes.

2689 **5.10.4.1 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.
```

2690 **5.10.4.2 Notes:**

- 2691
- None

2692 **5.10.5 tosca.policies.Performance**

2693 This is the default (root) TOSCA Policy Type definition that is used to declare performance requirements for
2694 TOSCA nodes or groups of nodes.

2695 **5.10.5.1 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.
```

2696 **5.10.5.2 Notes:**

- 2697
- None
- 2698

6 TOSCA Cloud Service Archive (CSAR) format

This section is **normative** and defines changes to the TOSCA archive format relative to the TOSCA v1.0 XML specification.

TOSCA Simple Profile definitions along with all accompanying artifacts (e.g. scripts, binaries, configuration files) can be packaged together in a CSAR file as already defined in the TOSCA version 1.0 specification [TOSCA-1.0]. In contrast to the TOSCA 1.0 CSAR file specification (see chapter 16 in [TOSCA-1.0]), this simple profile makes a few simplifications both in terms of overall CSAR file structure as well as meta-file content as described below.

6.1 Overall Structure of a CSAR

A CSAR zip file is required to contain a **TOSCA-Metadata** directory, which in turn contains the **TOSCA.meta** metadata file that provides entry information for a TOSCA orchestrator processing the CSAR file.

The CSAR file may contain other directories with arbitrary names and contents. Note that in contrast to the TOSCA 1.0 specification, it is not required to put TOSCA definitions files into a special “Definitions” directory, but definitions YAML files can be placed into any directory within the CSAR file.

6.2 TOSCA Meta File

The **TOSCA.meta** file structure follows the exact same syntax as defined in the TOSCA 1.0 specification. However, it is only required to include *block_0* (see section 16.2 in [TOSCA-1.0]) with the **Entry-Definitions** keyword pointing to a valid TOSCA definitions YAML file that a TOSCA orchestrator should use as entry for parsing the contents of the overall CSAR file.

Note that it is not required to explicitly list TOSCA definitions files in subsequent blocks of the **TOSCA.meta** file, but any TOSCA definitions files besides the one denoted by the **Entry-Definitions** keyword can be found by a TOSCA orchestrator by processing respective **imports** statements in the entry definitions file (or in recursively imported files).

Note also that any additional artifact files (e.g. scripts, binaries, configuration files) do not have to be declared explicitly through blocks in the **TOSCA.meta** file. Instead, such artifacts will be fully described and pointed to by relative path names through artifact definitions in one of the TOSCA definitions files contained in the CSAR.

Due to the simplified structure of the CSAR file and **TOSCA.meta** file compared to TOSCA 1.0, the **CSAR-Version** keyword listed in *block_0* of the meta-file is required to denote version **1.1**.

6.2.1 Example

The following listing represents a valid **TOSCA.meta** file according to this TOSCA Simple Profile specification.

```
TOSCA-Meta-File-Version: 1.0
CSAR-Version: 1.1
Created-By: OASIS TOSCA TC
Entry-Definitions: definitions/tosca_elk.yaml
```

This **TOSCA.meta** file indicates its simplified TOSCA Simple Profile structure by means of the **CSAR-Version** keyword with value **1.1**. The **Entry-Definitions** keyword points to a TOSCA definitions YAML file with the name **tosca_elk.yaml** which is contained in a directory called **definitions** within the root of the CSAR file.

7 TOSCA networking

This section is **normative** and describes how to express and control the application centric network semantics available in TOSCA.

7.1 Networking and Service Template Portability

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.

TOSCA Networking takes the following approach.

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.

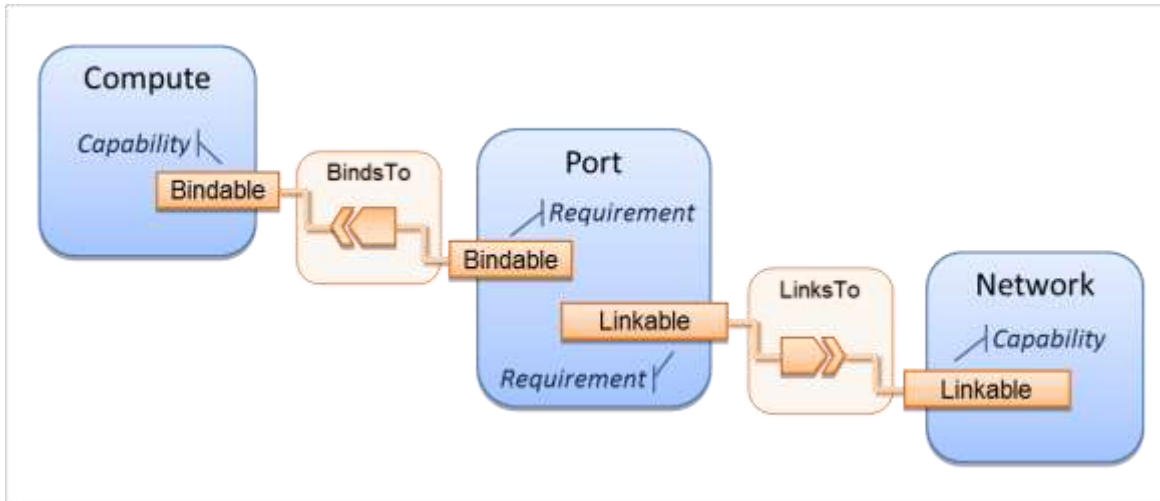
7.2 Connectivity Semantics

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.

Software components are expressed as Node Types in TOSCA which can express virtually any kind of concept in a TOSCA model. Node Types offering network based functions can model their connectivity using a special Endpoint Capability, [tosca.capabilities.Endpoint](#), designed for this purpose. Node Types which require an Endpoint can specify this as a TOSCA requirement. A special Relationship Type, [tosca.relationships.ConnectsTo](#), is used to implicitly or explicitly relate the source Node Type's endpoint to the required endpoint in the target node type. Since [tosca.capabilities.Endpoint](#) and [tosca.relationships.ConnectsTo](#) are TOSCA types, they can be used

2781 in templates and extended by subclassing in the usual ways, thus allowing the expression of additional semantics
2782 as needed.

2783 The following diagram shows how the TOSCA node, capability and relationship types enable modeling the
2784 application layer decoupled from the network model intersecting at the Compute node using the [Bindable](#)
2785 capability type.



2786 As you can see, the Port node type effectively acts a broker node between the Network node description and a
2787 host Compute node of an application.

2788 7.3 Expressing connectivity semantics

2789 This section describes how TOSCA supports the typical client/server and group communication semantics found
2790 in application architectures.

2791 7.3.1 Connection initiation semantics

2792 The `tosca.relationships.ConnectsTo` expresses that requirement that a source application component needs to be
2793 able to communicate with a target software component to consume the services of the target. `ConnectTo` is a
2794 component interdependency semantic in the most general sense and does not try imply how the communication
2795 between the source and target components is physically realized.

2796

2797 Application component intercommunication typically has conventions regarding which component(s) initiate the
2798 communication. Connection initiation semantics are specified in [tosca.capabilities.Endpoint](#). Endpoints at each
2799 end of the [tosca.relationships.ConnectsTo](#) must indicate identical connection initiation semantics.

2800

2801 The following sections describe the normative connection initiation semantics for the
2802 `tosca.relationships.ConnectsTo` Relationship Type.

2803 7.3.1.1 Source to Target

2804 The Source to Target communication initiation semantic is the most common case where the source component
2805 initiates communication with the target component in order to fulfill an instance of the
2806 `tosca.relationships.ConnectsTo` relationship. The typical case is a “client” component connecting to a “server”
2807 component where the client initiates a stream oriented connection to a pre-defined transport specific port or set of
2808 ports.

2809

2810 It is the responsibility of the TOSCA implementation to ensure the source component has a suitable network path
2811 to the target component and that the ports specified in the respective [tosca.capabilities.Endpoint](#) are not blocked.
2812 The TOSCA implementation may only represent state of the `tosca.relationships.ConnectsTo` relationship as

2813 fulfilled after the actual network communication is enabled and the source and target components are in their
2814 operational states.

2815

2816 Note that the connection initiation semantic only impacts the fulfillment of the actual connectivity and does not
2817 impact the node traversal order implied by the `tosca.relationships.ConnectsTo` Relationship Type.

2818 **7.3.1.2 Target to Source**

2819 The Target to Source communication initiation semantic is a less common case where the target component
2820 initiates communication with the source component in order to fulfill an instance of the
2821 `tosca.relationships.ConnectsTo` relationship. This “reverse” connection initiation direction is typically required due
2822 to some technical requirements of the components or protocols involved, such as the requirement that SSH must
2823 only be initiated from target component in order to fulfill the services required by the source component.

2824

2825 It is the responsibility of the TOSCA implementation to ensure the source component has a suitable network path
2826 to the target component and that the ports specified in the respective `tosca.capabilities.Endpoint` are not blocked.
2827 The TOSCA implementation may only represent state of the `tosca.relationships.ConnectsTo` relationship as
2828 fulfilled after the actual network communication is enabled and the source and target components are in their
2829 operational states.

2830

2831 Note that the connection initiation semantic only impacts the fulfillment of the actual connectivity and does not
2832 impact the node traversal order implied by the `tosca.relationships.ConnectsTo` Relationship Type.

2833

2834 **7.3.1.3 Peer-to-Peer**

2835 The Peer-to-Peer communication initiation semantic allows any member of a group to initiate communication with
2836 any other member of the same group at any time. This semantic typically appears in clustering and distributed
2837 services where there is redundancy of components or services.

2838

2839 It is the responsibility of the TOSCA implementation to ensure the source component has a suitable network path
2840 between all the member component instances and that the ports specified in the respective
2841 `tosca.capabilities.Endpoint` are not blocked, and the appropriate multicast communication, if necessary, enabled.
2842 The TOSCA implementation may only represent state of the `tosca.relationships.ConnectsTo` relationship as
2843 fulfilled after the actual network communication is enabled such that at least one member component of the group
2844 may reach any other member component of the group.

2845

2846 Endpoints specifying the Peer-to-Peer initiation semantic need not be related with a
2847 `tosca.relationships.ConnectsTo` relationship for the common case where the same set of component instances
2848 must communicate with each other.

2849

2850 Note that the connection initiation semantic only impacts the fulfillment of the actual connectivity and does not
2851 impact the node traversal order implied by the `tosca.relationships.ConnectsTo` Relationship Type.

2852 **7.3.2 Specifying layer 4 ports**

2853 TOSCA Service Templates must express enough details about application component intercommunication to
2854 enable TOSCA implementations to fulfill these communication semantics in the network infrastructure. TOSCA
2855 currently focuses on TCP/IP as this is the most pervasive in today's cloud infrastructures. The layer 4 ports
2856 required for application component intercommunication are specified in `tosca.capabilities.Endpoint`. The union of
2857 the port specifications of both the source and target `tosca.capabilities.Endpoint` which are part of the
2858 `tosca.relationships.ConnectsTo` Relationship Template are interpreted as the effective set of ports which must be
2859 allowed in the network communication.

2860

2861 The meaning of Source and Target port(s) corresponds to the direction of the respective
2862 tosca.relationships.ConnectsTo.

2863 **7.4 Network provisioning**

2864 **7.4.1 Declarative network provisioning**

2865 TOSCA orchestrators are responsible for the provisioning of the network connectivity for declarative TOCSA
2866 Service Templates (Declarative TOCSA Service Templates don't contain explicit plans). This means that the
2867 TOSCA orchestrator must be able to infer a suitable logical connectivity model from the Service Template and
2868 then decide how to provision the logical connectivity, referred to as "fulfillment", on the available underlying
2869 infrastructure. In order to enable fulfillment, sufficient technical details still must be specified, such as the required
2870 protocols, ports and QOS information. TOSCA connectivity types, such as tosca.capabilities.Endpoint, provide
2871 well defined means to express these details.

2872 **7.4.2 Implicit network fulfillment**

2873 TOSCA Service Templates are by default network agnostic. TOSCA's application centric approach only requires
2874 that a TOSCA Service Template contain enough information for a TOSCA orchestrator to infer suitable network
2875 connectivity to meet the needs of the application components. Thus Service Template designers are not required
2876 to be aware of or provide specific requirements for underlying networks. This approach yields the most portable
2877 Service Templates, allowing them to be deployed into any infrastructure which can provide the necessary
2878 component interconnectivity.

2879 **7.4.3 Controlling network fulfillment**

2880 TOSCA provides mechanisms for providing control over network fulfillment.
2881 This mechanism allows the application network designer to express in service template or network template how
2882 the networks should be provisioned.

2883
2884 For the use cases described below let's assume we have a typical 3-tier application which is consisting of FE
2885 (frontend), BE (backend) and DB (database) tiers. The simple application topology diagram can be shown below:




Figure-5: Typical 3-Tier Network

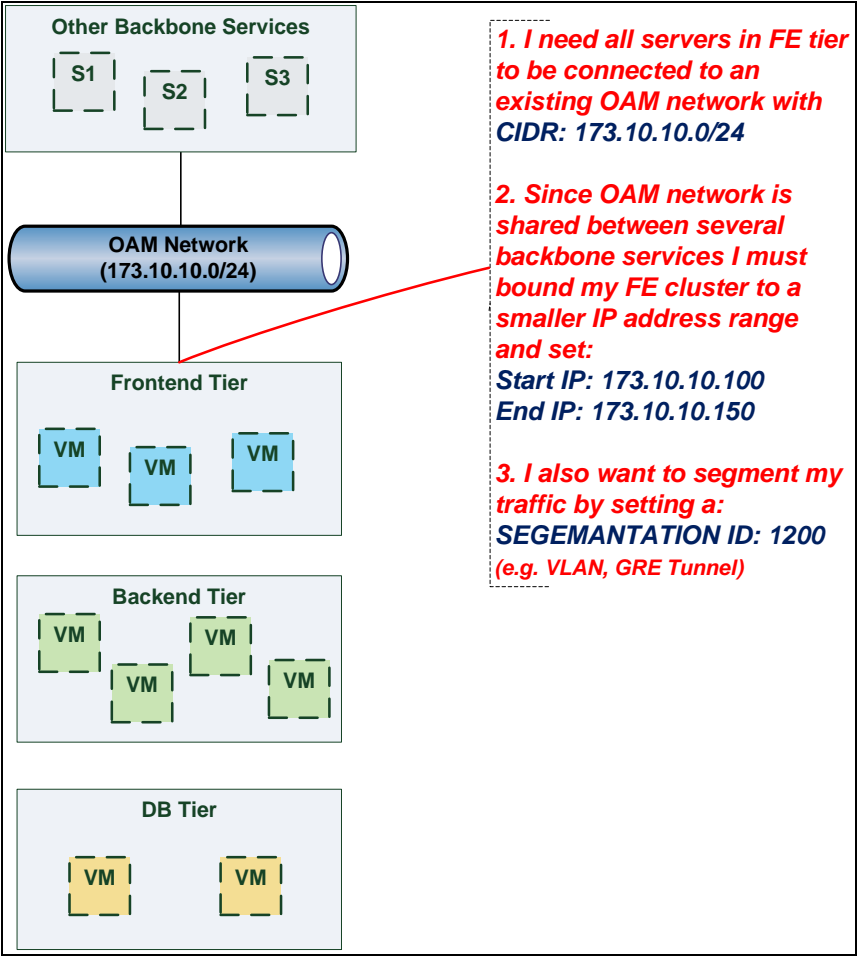
2888 **7.4.3.1 Use case: OAM Network**

2889 When deploying an application in service provider's on-premise cloud, it's very common that one or more of the
2890 application's services should be accessible from an ad-hoc OAM (Operations, Administration and Management)
2891 network which exists in the service provider backbone.

2892
2893 As an application network designer, I'd like to express in my TOSCA network template (which corresponds to my
2894 TOSCA service template) the network CIDR block, start ip, end ip and segmentation ID (e.g. VLAN id).

2895 The diagram below depicts a typical 3-tiers application with specific networking requirements for its FE tier server
2896 cluster:


2897



2898

2899 **7.4.3.2 Use case: Data Traffic network**

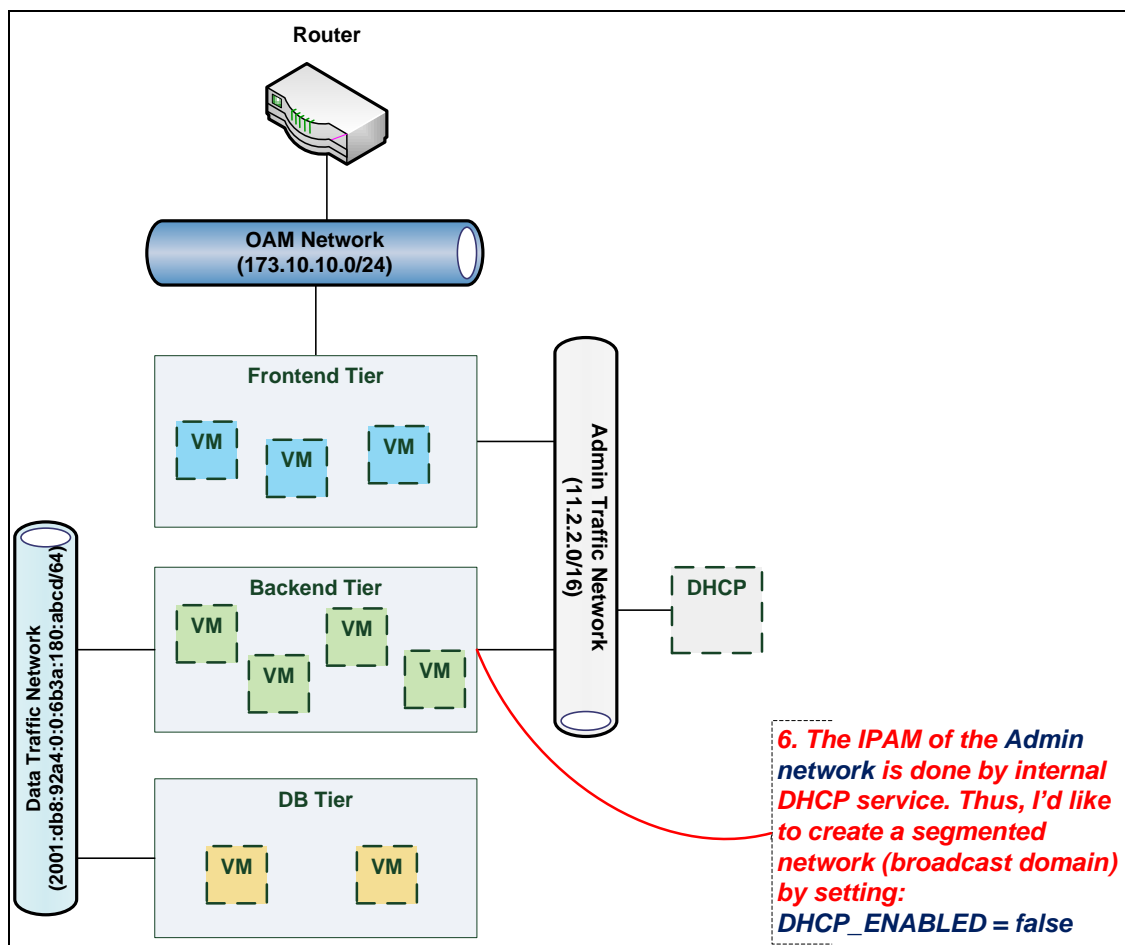
2900 The diagram below defines a set of networking requirements for the backend and DB tiers of the 3-tier app
2901 mentioned above.



7.4.3.3 Use case: Bring my own DHCP

The same 3-tier app requires for its admin traffic network to manage the IP allocation by its own DHCP which runs autonomously as part of application domain.

For this purpose, the app network designer would like to express in TOSCA that the underlying provisioned network will be set with DHCP_ENABLED=false. See this illustrated in the figure below:



7.5 Network Types

7.5.1 toska.nodes.network.Network

The TOSCA **Network** node represents a simple, logical network service.

Shorthand Name	Network
Type Qualified Name	tosca:Network
Type URI	tosca.nodes.network.Network

7.5.1.1 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
gateway_ip	no	string	None	The gateway IP address.

Name	Required	Type	Constraints	Description
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.

2914 7.5.1.2 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.

2915 7.5.1.3 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

```

2916 7.5.1.4 Additional Requirements

- 2917 • None

2918 7.5.2 **tosca.nodes.network.Port**

2919 The TOSCA **Port** node represents a logical entity that associates between Compute and Network normative
2920 types.

2921 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

2922 7.5.2.1 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 fulfil 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.

2923 7.5.2.2 Attributes

Name	Required	Type	Constraints	Description
ip_address	no	string	None	The IP address would be assigned to the associated compute instance.

2924 7.5.2.3 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.LinkTo
    - binding:
        capability: tosca.capabilities.network.Bindable
        relationship: tosca.relationships.network.BindsTo

```

7.5.2.4 Additional Requirements

- None

7.5.3 [tosca.capabilities.network.Linkable](#)

A node type that includes the Linkable capability indicates that it can be pointed by [tosca.relationships.network.LinkTo](#) relationship type.

Shorthand Name	Linkable
Type Qualified Name	tosca::Linkable
Type URI	tosca.capabilities.network.Linkable

7.5.3.1 Properties

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

7.5.3.2 Definition

```

tosca.capabilities.network.Linkable:
  derived_from: tosca.capabilities.Node

```

7.5.4 [tosca.relationships.network.LinkTo](#)

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

7.5.4.1 Definition

```
tosca.relationships.network.LinksTo:
  derived_from: tosca.relationships.DependsOn
  valid_target_types: [ tosca.capabilities.network.Linkable ]
```

7.5.5 tosca.relationships.network.BindsTo

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

7.5.5.1 Definition

```
tosca.relationships.network.BindsTo:
  derived_from: tosca.relationships.DependsOn
  valid_target_types: [ tosca.capabilities.network.Bindable ]
```

7.6 Network modeling approaches

7.6.1 Option 1: Specifying a network outside the application's Service Template

This approach allows someone who understands the application's networking requirements, mapping the details of the underlying network to the appropriate node templates in the application.

The motivation for this approach is providing the application network designer a fine-grained control on how networks are provisioned and stitched to its application by the TOSCA orchestrator and underlying cloud infrastructure while still preserving the portability of his service template. Preserving the portability means here not doing any modification in service template but just "plug-in" the desired network modeling. The network modeling can reside in the same service template file but the best practice should be placing it in a separated self-contained network template file.

This "pluggable" network template approach introduces a new normative node type called Port, capability called [tosca.capabilities.network.Linkable](#) and relationship type called [tosca.relationships.network.LinksTo](#).

The idea of the Port is to elegantly associate the desired compute nodes with the desired network nodes while not "touching" the compute itself.

The following diagram series demonstrate the plug-ability strength of this approach.

Let's assume an application designer has modeled a service template as shown in Figure 1 that describes the application topology nodes (compute, storage, software components, etc.) with their relationships. The designer ideally wants to preserve this service template and use it in any cloud provider environment without any change.




Figure-6: Generic Service Template

When the application designer comes to consider its application networking requirement they typically call the network architect/designer from their company (who has the correct expertise).

The network designer, after understanding the application connectivity requirements and optionally the target cloud provider environment, is able to model the network template and plug it to the service template as shown in Figure 2:

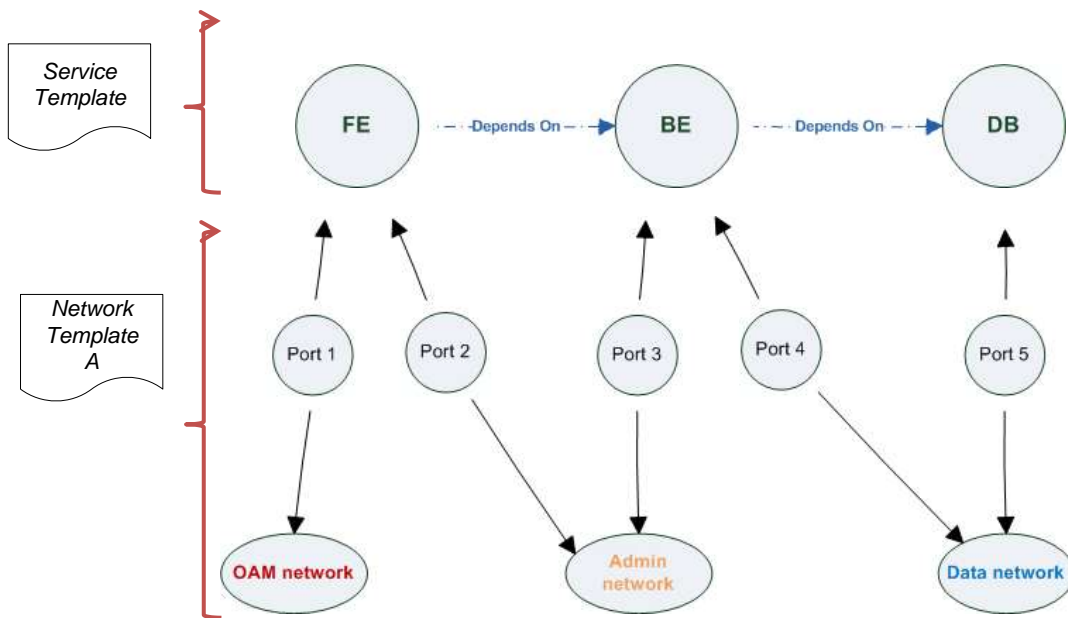


Figure-7: Service template with network template A

When there's a new target cloud environment to run the application on, the network designer is simply creates a new network template B that corresponds to the new environmental conditions and provide it to the application designer which packs it into the application CSAR.

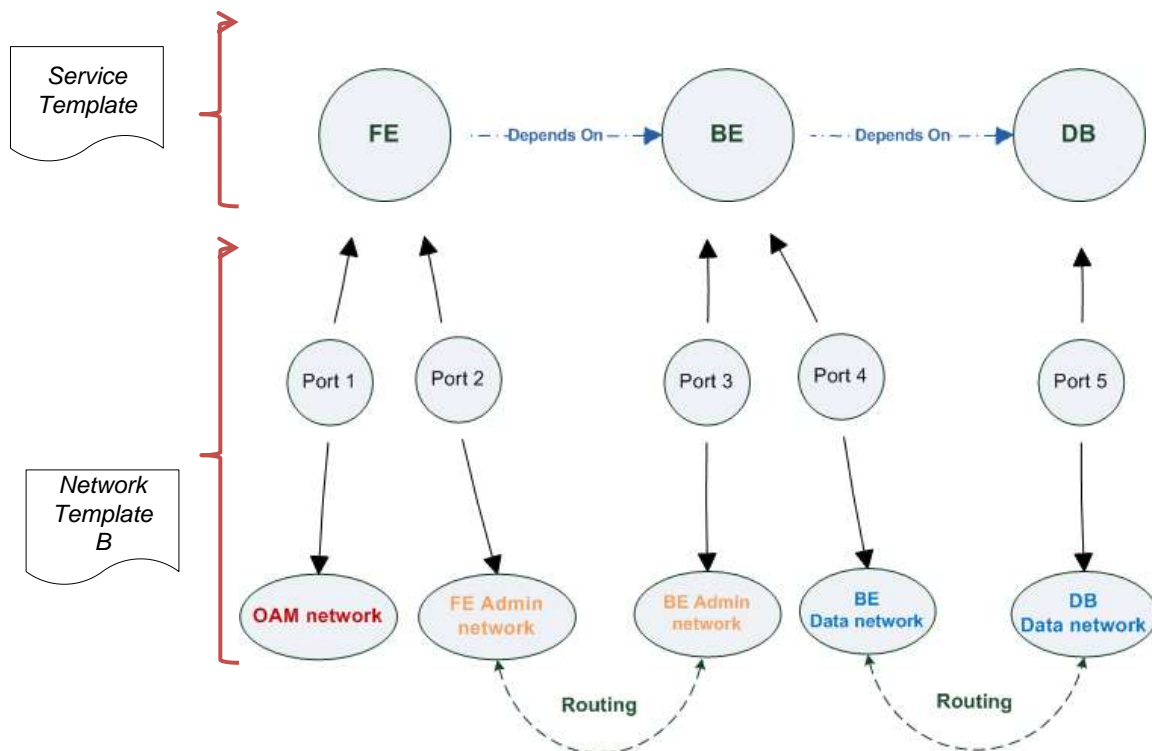


Figure-8: Service template with network template B

The node templates for these three networks would be defined as follows:

```
node_templates:
  frontend:
    type: toska.nodes.Compute
    properties: # omitted for brevity

  backend:
    type: toska.nodes.Compute
    properties: # omitted for brevity

  database:
    type: toska.nodes.Compute
    properties: # omitted for brevity

  oam_network:
    type: toska.nodes.network.Network
    properties: # omitted for brevity

  admin_network:
    type: toska.nodes.network.Network
    properties: # omitted for brevity

  data_network:
```

```

    type: toska.nodes.network.Network
    properties: # omitted for brevity

# ports definition
fe_oam_net_port:
  type: toska.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: toska.nodes.network.Port
  requirements:
    - link: admin_network
    - binding: frontend

be_admin_net_port:
  type: toska.nodes.network.Port
  properties:
    order: 0
  requirements:
    - link: admin_network
    - binding: backend

be_data_net_port:
  type: toska.nodes.network.Port
  properties:
    order: 1
  requirements:
    - link: data_network
    - binding: backend

db_data_net_port:
  type: toska.nodes.network.Port
  requirements:
    - link: data_network
    - binding: database

```

2974
2975
2976
2977
2978

7.6.2 Option 2: Specifying network requirements within the application's Service Template

This approach allows the Service Template designer to map an endpoint to a logical network.

The use case shown below examines a way to express in the TOSCA YAML service template a typical 3-tier application with their required networking modeling:

```
node_templates:
  frontend:
    type: toska.nodes.Compute
    properties: # omitted for brevity
    requirements:
      - network_oam: oam_network
      - network_admin: admin_network
  backend:
    type: toska.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 }
```

2979

8 Non-normative type definitions

This section defines **non-normative** types used in examples or use cases within this specification.

8.1 Artifact Types

This section contains are non-normative Artifact Types used in use cases and examples.

8.1.1 `tosca.artifacts.Deployment.Image.Container.Docker`

This artifact represents a Docker “image” (a TOSCA deployment artifact type) which is a binary comprised of one or more (a union of read-only and read-write) layers created from snapshots within the underlying Docker **Union File System**.

8.1.1.1 Definition

```
tosca.artifacts.Deployment.Image.Container.Docker:
  derived_from: tosca.artifacts.Deployment.Image
  description: Docker Container Image
```

8.1.2 `tosca.artifacts.Deployment.Image.VM.ISO`

A Virtual Machine (VM) formatted as an ISO standard disk image.

8.1.2.1 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 ]
```

8.1.3 `tosca.artifacts.Deployment.Image.VM.QCOW2`

A Virtual Machine (VM) formatted as a QEMU emulator version 2 standard disk image.

8.1.3.1 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 ]
```

8.2 Capability Types

This section contains are non-normative Capability Types used in use cases and examples.

8.2.1 `tosca.capabilities.Container.Docker`

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

2999

8.2.1.1 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.

3000

8.2.1.2 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
```

8.2.1.3 Additional requirements

- When the **expose_ports** property is used, only the **source** and **source_range** properties of **PortSpec** **SHALL** be valid for supplying port numbers or ranges, the **target** and **target_range** properties are ignored.

8.3 Node Types

This section contains non-normative node types referenced in use cases and examples. All additional Attributes, Properties, Requirements and Capabilities shown in their definitions (and are not inherited from ancestor normative types) are also considered to be non-normative.

8.3.1 **tosca.nodes.Database.MySQL**

8.3.1.1 Properties

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

8.3.1.2 Definition

```
tosca.nodes.Database.MySQL:
  derived_from: toska.nodes.Database
  requirements:
    - host:
        node: toska.nodes.DBMS.MySQL
```

8.3.2 **tosca.nodes.DBMS.MySQL**

8.3.2.1 Properties

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

8.3.2.2 Definition

```
tosca.nodes.DBMS.MySQL:
  derived_from: toska.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 ]

```

3015 8.3.3 [tosca.nodes.WebServer.Apache](#)

3016 8.3.3.1 Properties

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

3017 8.3.3.2 Definition

```

tosca.nodes.WebServer.Apache:
    derived_from: tosca.nodes.WebServer

```

3018 8.3.4 [tosca.nodes.WebApplication.WordPress](#)

3019 8.3.4.1 Properties

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

3020 8.3.4.2 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

```

3021 **8.3.5 `tosca.nodes.WebServer.Nodejs`**

3022 **8.3.5.1 Properties**

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

3023 **8.3.5.2 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
```

3024 **8.3.6 `tosca.nodes.Container.Application.Docker`**

3025 **8.3.6.1 Properties**

Name	Required	Type	Constraints	Description
N/A	N/A	N/A	N/A	N/A

3026 **8.3.6.2 Definition**

```
tosca.nodes.Container.Application.Docker:
  derived_from: tosca.nodes.Container.Application
  requirements:
    - host:
        capability: tosca.capabilities.Container.Docker
```

9 Component Modeling Use Cases

This section is **non-normative** and includes use cases that explore how to model components and their relationships using TOSCA Simple Profile in YAML.

9.1.1 Use Case: Exploring the HostedOn relationship using WebApplication and WebServer

This use case examines the ways TOSCA YAML can be used to express a simple hosting relationship (i.e., **HostedOn**) using the normative TOSCA **WebServer** and **WebApplication** node types defined in this specification.

9.1.1.1 WebServer declares its “host” capability

For convenience, relevant parts of the normative TOSCA Node Type for **WebServer** are shown below:

```
tosca.nodes.WebServer
  derived_from: SoftwareComponent
  capabilities:
    ...
  host:
    type: tosca.capabilities.Container
    valid_source_types: [ tosca.nodes.WebApplication ]
```

As can be seen, the **WebServer** Node Type declares its capability to “contain” (i.e., host) other nodes using the symbolic name “**host**” and providing the Capability Type **tosca.capabilities.Container**. It should be noted that the symbolic name of “**host**” is not a reserved word, but one assigned by the type designer that implies at or betokens the associated capability. The **Container** capability definition also includes a required list of valid Node Types that can be contained by this, the **WebServer**, Node Type. This list is declared using the keyname of **valid_source_types** and in this case it includes only allowed type **WebApplication**.

9.1.1.2 WebApplication declares its “host” requirement

The **WebApplication** node type needs to be able to describe the type of capability a target node would have to provide in order to “host” it. The normative TOSCA capability type **tosca.capabilities.Container** is used to describe all normative TOSCA hosting (i.e., container-containee pattern) relationships. As can be seen below, the **WebApplication** accomplishes this by declaring a requirement with the symbolic name “**host**” with the **capability** keyname set to **tosca.capabilities.Container**.

Again, for convenience, the relevant parts of the normative **WebApplication** Node Type are shown below:

```
tosca.nodes.WebApplication:
  derived_from: tosca.nodes.Root
  requirements:
    - host:
        capability: tosca.capabilities.Container
        node: tosca.nodes.WebServer
        relationship: tosca.relationships.HostedOn
```

3049 9.1.1.2.1 Notes

- 3050 • The symbolic name “host” is not a keyword and was selected for consistent use in TOSCA normative
3051 node types to give the reader an indication of the type of requirement being referenced. A valid
3052 HostedOn relationship could still be established between WebApplicaton and WebServer in a TOSCA
3053 Service Template regardless of the symbolic name assigned to either the requirement or capability
3054 declaration.

3055 9.1.2 Use Case: Establishing a ConnectsTo relationship to WebServer

3056 This use case examines the ways TOSCA YAML can be used to express a simple connection relationship (i.e.,
3057 [ConnectsTo](#)) between some service derived from the [SoftwareComponent](#) Node Type, to the normative
3058 [WebServer](#) node type defined in this specification.

3059 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
```

3060 Since the normative **WebServer** Node Type only declares one capability of type **tosca.capabilities.Endpoint**
3061 (or **Endpoint**, its shortname alias in TOSCA) using the symbolic name **data_endpoint**, the **my_web_service**
3062 node template does not need to declare that symbolic name on its requirement declaration. If however, the
3063 **my_web_server** node was based upon some other node type that declared more than one capability of type
3064 **Endpoint**, then the **capability** keyname could be used to supply the desired symbolic name if necessary.

9.1.2.1 Best practice

It should be noted that the best practice for designing Node Types in TOSCA should not export two capabilities of the same type if they truly offer different functionality (i.e., different capabilities) which should be distinguished using different Capability Type definitions.

9.1.3 Use Case: Attaching (local) BlockStorage to a Compute node

This use case examines the ways TOSCA YAML can be used to express a simple AttachesTo relationship between a Compute node and a locally attached BlockStorage node.

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
```

9.1.4 Use Case: Reusing a BlockStorage Relationship using Relationship Type or Relationship Template

This builds upon the previous use case (9.1.3) to examine how a template author could attach multiple Compute nodes (templates) to the same BlockStorage node (template), but with slightly different property values for the AttachesTo relationship.

Specifically, several notation options are shown (in this use case) that achieve the same desired result.

9.1.4.1 Simple Profile Rationale

Referencing an explicitly declared Relationship Template is a convenience of the Simple Profile that allows template authors an entity to set, constrain or override the properties and operations as defined in its declared (Relationship) Type much as allowed now for Node Templates. It is especially useful when a complex Relationship Type (with many configurable properties or operations) has several logical occurrences in the same Service (Topology) Template; allowing the author to avoid configuring these same properties and operations in multiple Node Templates.

9.1.4.2 Notation Style #1: Augment AttachesTo Relationship Type directly in each Node Template

This notation extends the methodology used for establishing a HostedOn relationship, but allowing template author to supply (dynamic) configuration and/or override of properties and operations.

Note: This option will remain valid for Simple Profile regardless of other notation (copy or aliasing) options being discussed or adopted for future versions.

```
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
```

3095

3096

3097

9.1.4.3 Notation Style #2: Use the ‘template’ keyword on the Node Templates to specify which named Relationship Template to use

3098

3099

3100

This option shows how to explicitly declare different named Relationship Templates within the Service Template as part of a **relationship_templates** section (which have different property values) and can be referenced by different Compute typed Node Templates.

3101

```
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

```

3102

3103 9.1.4.4 Notation Style #3: Using the “copy” keyname to define a similar Relationship 3104 Template

3105 How does TOSCA make it easier to create a new relationship template that is mostly the same as one that exists
3106 without manually copying all the same information? TOSCA provides the **copy** keyname as a convenient way to
3107 copy an existing template definition into a new template definition as a starting point or basis for describing a new
3108 definition and avoid manual copy. The end results are cleaner TOSCA Service Templates that allows the
3109 description of only the changes (or deltas) between similar templates.

3110 The example below shows that the Relationship Template named **storage_attachesto_1** provides some
3111 overrides (conceptually a large set of overrides) on its Type which the Relationship Template named
3112 **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
```

10 Application Modeling Use Cases

This section is **non-normative** and includes use cases that show how to model Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS) and complete application use cases using TOSCA Simple Profile in YAML.

10.1 Use cases

10.1.1 Overview

Use Case		
Name	Description	Service Template link (Entry Definition)
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.	None
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.	None
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.	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/storage/tosca_blockstorage_with_attachment.yaml#L19
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.	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/storage/tosca_blockstorage_with_custom_relationship_type.yaml
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.	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/storage/tosca_blockstorage_with_relationship_template.yaml
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.	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/storage/tosca_blockstorage_with_attachment_notation1.yaml
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.	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/storage/tosca_blockstorage_with_attachment_notation2.yaml
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	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/storage/tosca_multiple_blockstorage_with_attachment.yaml

Use Case		
Name	Description	Service Template link (Entry Definition)
	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.	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/storage/tosca_single_object_store.yaml
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.	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/network/tosca_one_server_one_network.yaml
Network-2: Server bound to an existing network	Shows how to use a network_name as an input parameter to the template to allow a server to be associated with (i.e. bound to) an existing Network .	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/network/tosca_server_on_existing_network.yaml
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 .	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/network/tosca_two_servers_one_network.yaml
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).	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/network/tosca_one_server_three_networks.yaml
WebServer-DBMS-1: 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.	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/tosca_single_instance_wordpress.yaml
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.	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/tosca_nodejs_mongodb_two_instances.yaml
Multi-Tier-1: Elasticsearch, Logstash, Kibana (ELK)	<p>Shows Elasticsearch, Logstash and Kibana (ELK) being used in a typical manner to collect, search and monitor/visualize data from a running application.</p> <p>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.</p> <p>In addition to the application tiers, a 3rd tier is introduced with Logstash to collect data from the application tiers. Finally a 4th tier is added to search the Logstash data with Elasticsearch and visualize it using Kibana.</p> <p>Note: This use case also shows the convenience of using a single YAML macro (declared in the</p>	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/tosca_elk.yaml

Use Case		
Name	Description	Service Template link (Entry Definition)
	dsl_definitions section of the TOSCA Service Template) on multiple Compute nodes.	
Container-1: Containers using Docker single Compute instance (Containers only)	<p>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.</p> <p>Variation 1: Docker Container nodes (only) providing their Docker Requirements allowing platform (orchestrator) to select/provide the underlying Docker implementation (Capability).</p>	None

10.1.2 Compute: Create a single Compute instance with a host Operating System

10.1.2.1 Description

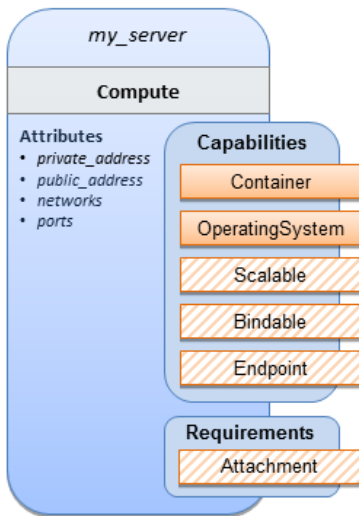
This use case demonstrates how the TOSCA Simple Profile specification can be used to stand up a single Compute instance with a guest Operating System using a normative TOSCA **Compute** node. The TOSCA Compute node is declarative in that the service template describes both the processor and host operating system platform characteristics (i.e., properties declared on the capability named “**os**” sometimes called a “flavor”) that are desired by the template author. The cloud provider would attempt to fulfill these properties (to the best of its abilities) during orchestration.

10.1.2.2 Features

This use case introduces the following TOSCA Simple Profile features:

- A node template that uses the normative TOSCA **Compute** Node Type along with showing an exemplary set of its properties being configured.
- Use of the TOSCA Service Template **inputs** section to declare a configurable value the template user may supply at runtime. In this case, the “**host**” property named “**num_cpus**” (of type integer) is declared.
 - Use of a property constraint to limit the allowed integer values for the “**num_cpus**” property to a specific list supplied in the property declaration.
- Use of the TOSCA Service Template **outputs** section to declare a value the template user may request at runtime. In this case, the property named “**instance_ip**” is declared
 - The “**instance_ip**” output property is programmatically retrieved from the **Compute** node’s “**public_address**” attribute using the TOSCA Service Template-level **get_attribute** function.

3138 10.1.2.3 Logical Diagram



3139

3140 10.1.2.4 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] }

```

10.1.2.5 Notes

- This use case uses a versioned, Linux Ubuntu distribution on the Compute node.

10.1.3 Software Component 1: Automatic deployment of a Virtual Machine (VM) image artifact

10.1.3.1 Description

This use case demonstrates how the TOSCA SoftwareComponent node type can be used to declare software that is packaged in a standard Virtual Machine (VM) image file format (i.e., in this case QCOW2) and is hosted on a TOSCA Compute node (instance). In this variation, the SoftwareComponent declares a VM image as a deployment artifact that 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.

10.1.3.2 Features

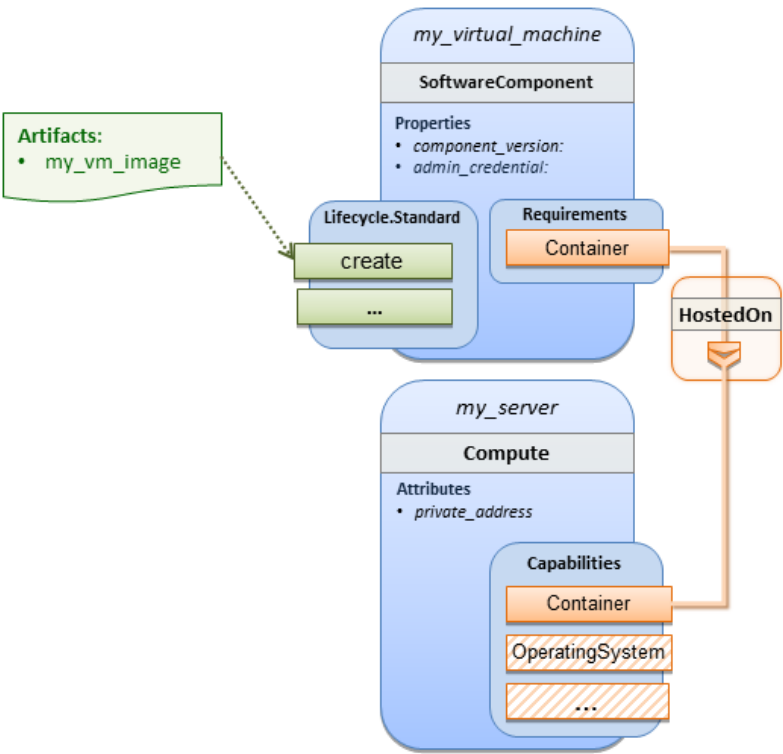
This use case introduces the following TOSCA Simple Profile features:

- A node template that uses the normative TOSCA **SoftwareComponent** Node Type along with showing an exemplary set of its properties being configured.
- Use of the TOSCA Service Template **artifacts** section to declare a Virtual Machine (VM) image artifact type which is referenced by the **SoftwareComponent** node template.
- The VM file format, in this case QCOW2, includes its own guest Operating System (OS) and therefore does **not** “require” a TOSCA **OperatingSystem** capability from the TOSCA Compute node.

10.1.3.3 Assumptions

This use case assumes the following:

- That the TOSCA Orchestrator (working with the Cloud provider’s underlying management services) is able to instantiate a Compute node that has a hypervisor that supports the Virtual Machine (VM) image format, in this case QCOW2, which should be compatible with many standard hypervisors such as XEN and KVM.
- This is not a “baremetal” use case and assumes the existence of a hypervisor on the machine that is allocated to “host” the Compute instance supports (e.g. has drivers, etc.) the VM image format in this example.



3170

3171 10.1.3.5 Sample YAML

```
tosca_definitions_version: tosca_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: tosca.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] }

```

10.1.3.6 Notes

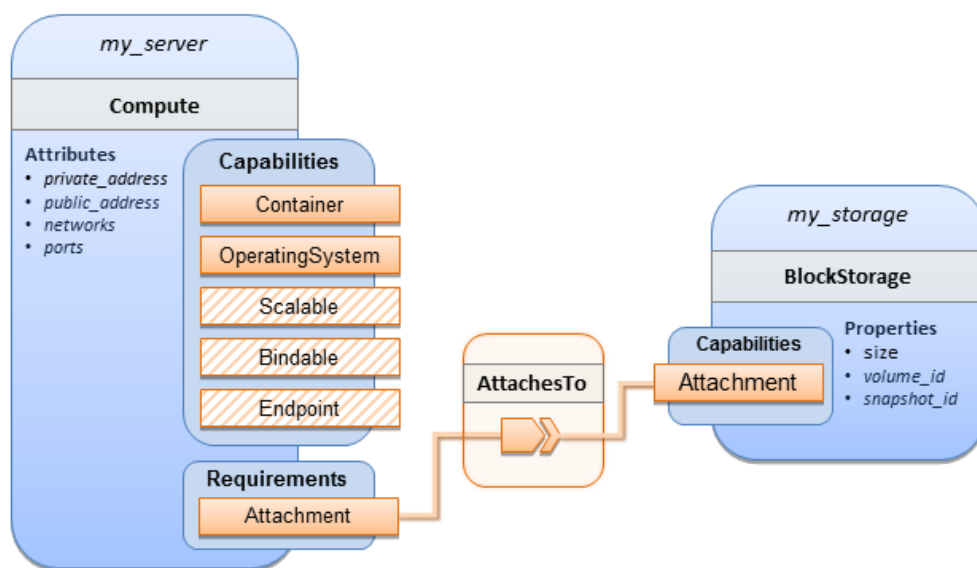
- The use of the **type** keyname on the **artifact** definition (within the **my_virtual_machine** node template) to declare the ISO image deployment artifact type (i.e., **tosca.artifacts.Deployment.Image.VM.ISO**) is redundant since the file extension is “.iso” which associated with this known, declared artifact type.
- This use case references a filename on the **my_vm_image** artifact, which indicates a Linux, Fedora 18, x86 VM image, only as one possible example.

10.1.4 Block Storage 1: Using the normative AttachesTo Relationship Type

10.1.4.1 Description

This use case demonstrates how to attach a TOSCA **BlockStorage** node to a **Compute** node using the normative **AttachesTo** relationship.

3183 10.1.4.2 Logical Diagram



3184

3185 10.1.4.3 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] }

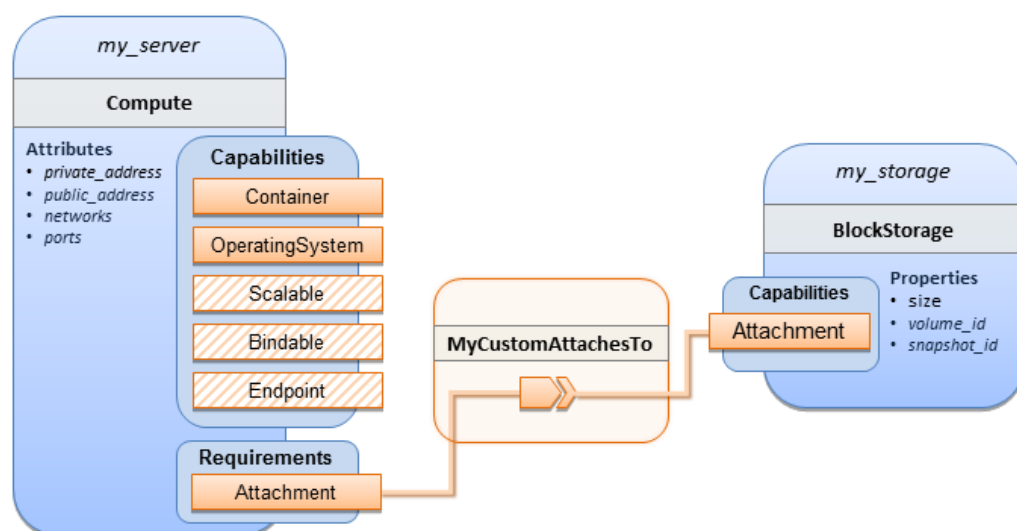
```

3186 10.1.5 Block Storage 2: Using a custom AttachesTo Relationship Type

3187 10.1.5.1 Description

3188 This use case demonstrates how to attach a TOSCA **BlockStorage** node to a **Compute** node using a custom
 3189 RelationshipType that derives from the normative **AttachesTo** relationship.

3190 10.1.5.2 Logical Diagram



3191

3192 10.1.5.3 Sample YAML

3193

```

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] }

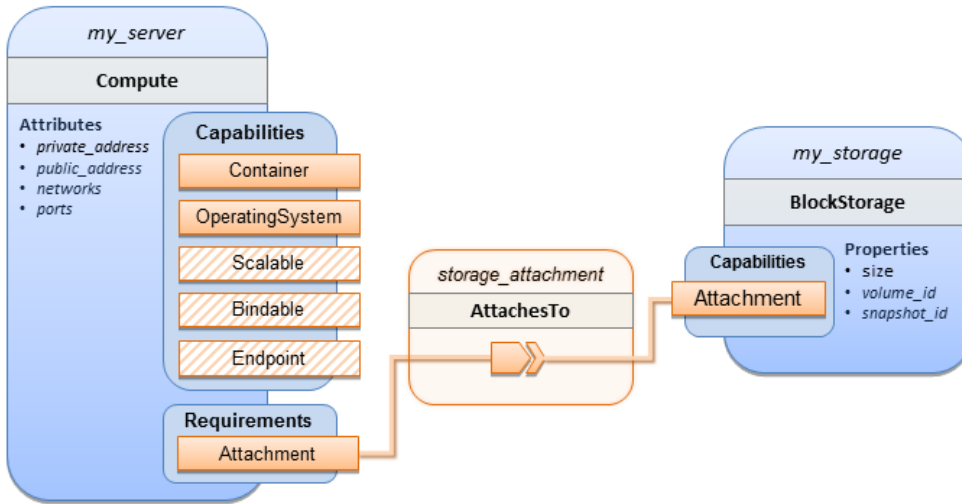
```

3194 10.1.6 Block Storage 3: Using a Relationship Template of type AttachesTo

3195 10.1.6.1 Description

3196 This use case demonstrates how to attach a TOSCA **BlockStorage** node to a **Compute** node using a TOSCA
3197 Relationship Template that is based upon the normative **AttachesTo** Relationship Type.

3198 10.1.6.2 Logical Diagram



3199

3200 10.1.6.3 Sample YAML

3201

```
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] }

```

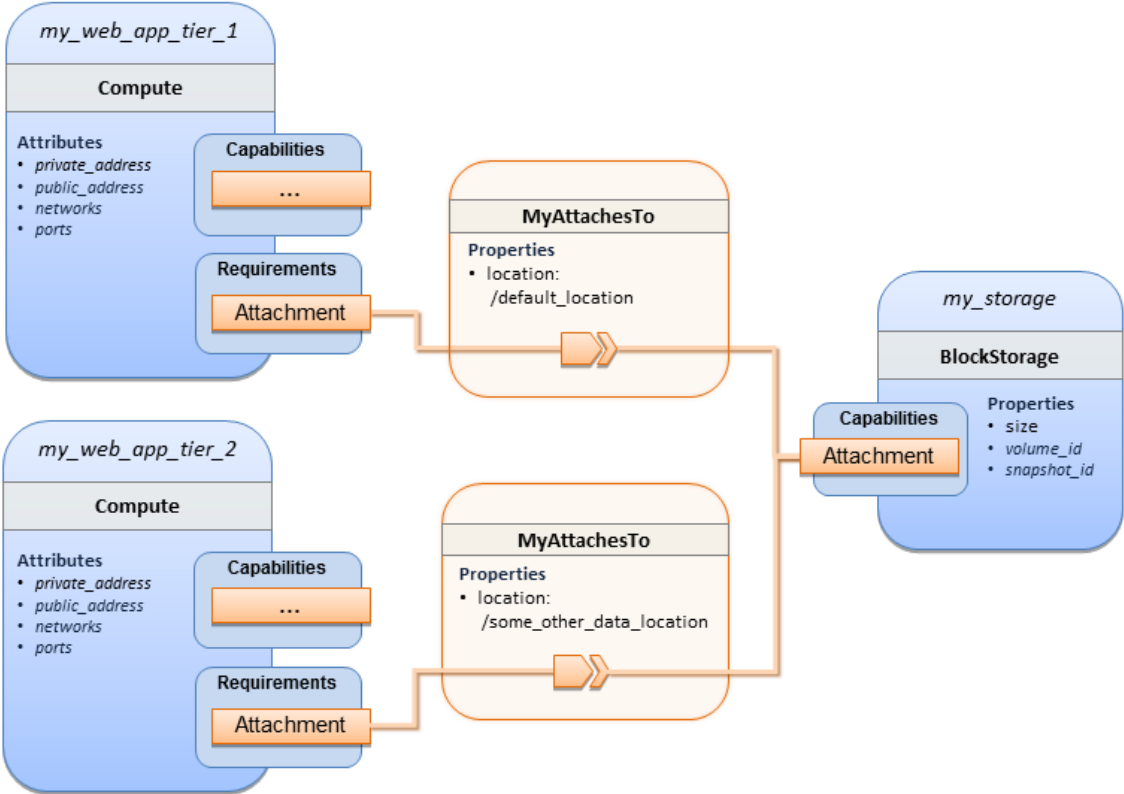

3202 **10.1.7 Block Storage 4: Single Block Storage shared by 2-Tier Application with**
3203 **custom AttachesTo Type and implied relationships**

3204 **10.1.7.1 Description**

3205 This use case shows 2 compute instances (2 tiers) with one BlockStorage node, and also uses a custom
3206 **AttachesTo** Relationship that provides a default mount point (i.e., **location**) which the 1st tier uses, but the 2nd
3207 tier provides a different mount point.

3208
3209 Please note that this use case assumes both Compute nodes are accessing different directories within the
3210 shared, block storage node to avoid collisions.

3211 **10.1.7.2 Logical Diagram**



3212
3213 **10.1.7.3 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] }

```


10.1.8 Block Storage 5: Single Block Storage shared by 2-Tier Application with custom AttachesTo Type and explicit Relationship Templates

10.1.8.1 Description

This use case is like the Notation1 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.

3221 Please note that this use case assumes both Compute nodes are accessing different directories within the
3222 shared, block storage node to avoid collisions.

3223 10.1.8.2 Logical Diagram



3224

3225 10.1.8.3 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: tosa.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: tosa.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: toska.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

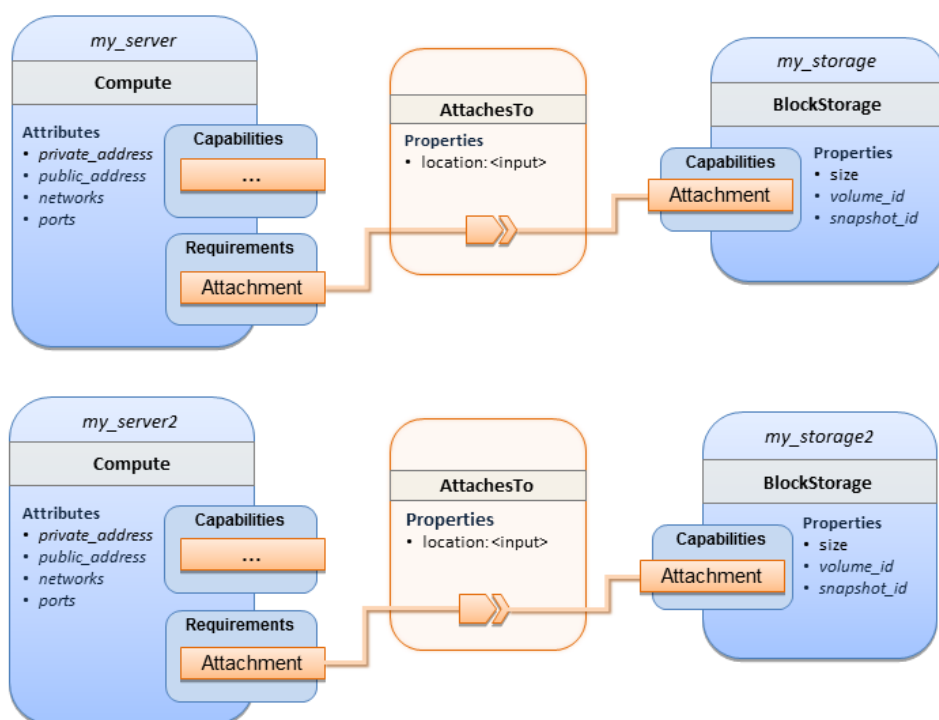
    outputs:
      private_ip_1:
        description: The private IP address of the application's first tier.
        value: { get_attribute: [my_web_app_tier_1, private_address] }
      private_ip_2:
        description: The private IP address of the application's second tier.
        value: { get_attribute: [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] }

```

3226 10.1.9 Block Storage 6: Multiple Block Storage attached to different Servers

3227 10.1.9.1 Description

3228 This use case demonstrates how two different TOSCA **BlockStorage** nodes can be attached to two different
 3229 **Compute** nodes (i.e., servers) each using the normative **AttachesTo** relationship.



3231

3232 10.1.9.3 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: tosa.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: tosa.nodes.BlockStorage
      properties:
        size: { get_input: storage_size }
        snapshot_id: { get_input: storage_snapshot_id }

  my_server2:
    type: tosa.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: tosa.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] }

```

3233 10.1.10 Object Storage 1: Creating an Object Storage service

3234 10.1.10.1 Description

3235 10.1.10.2 Logical Diagram



3236

3237 **10.1.10.3 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
```

3238 **10.1.11 Network 1: Server bound to a new network**

3239 **10.1.11.1 Description**

3240 Introduces the TOSCA **Network** and **Port** nodes used for modeling logical networks using the **LinksTo** and **BindsTo**
3241 Relationship Types. In this use case, the template is invoked without an existing network_name as an input
3242 property so a new network is created using the properties declared in the Network node.

3243 **10.1.11.2 Logical Diagram**



3244

```
tosca_definitions_version: tosa_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: tosa.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: tosa.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: tosa.nodes.network.Port
      requirements:
        - binding: my_server
```

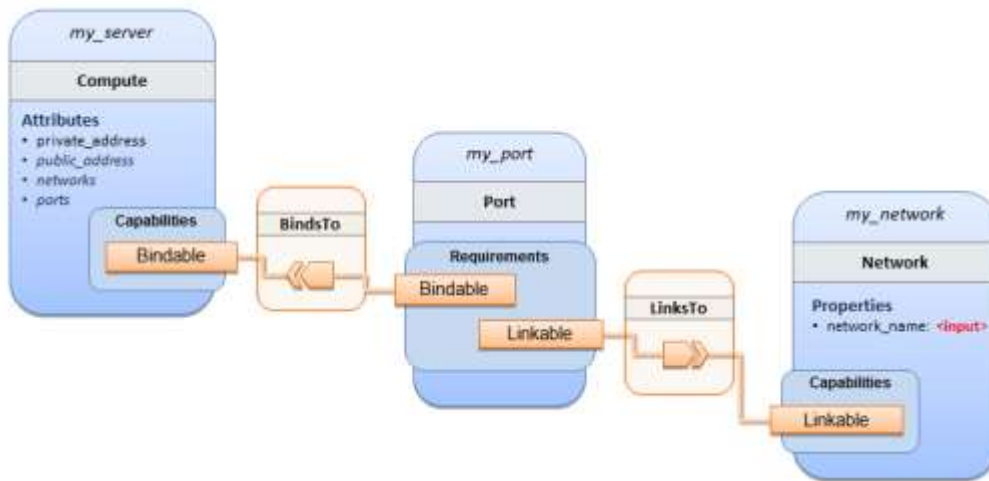
- link: my_network

10.1.12 Network 2: Server bound to an existing network

10.1.12.1 Description

This use case shows how to use a `network_name` as an input parameter to the template to allow a server to be associated with an existing network.

10.1.12.2 Logical Diagram



10.1.12.3 Sample YAML

```
tosca_definitions_version: tosca_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: toska.nodes.network.Network
  properties:
    network_name: { get_input: network_name }

my_port:
  type: toska.nodes.network.Port
  requirements:
    - binding:
        node: my_server
    - link:
        node: my_network


```

10.1.13 Network 3: Two servers bound to a single network

10.1.13.1 Description

This use case shows how two servers (**Compute** nodes) can be bound to the same **Network** (node) using two logical network **Ports**.

10.1.13.2 Logical Diagram



```
tosca_definitions_version: tosca_simple_yaml_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: 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:
    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

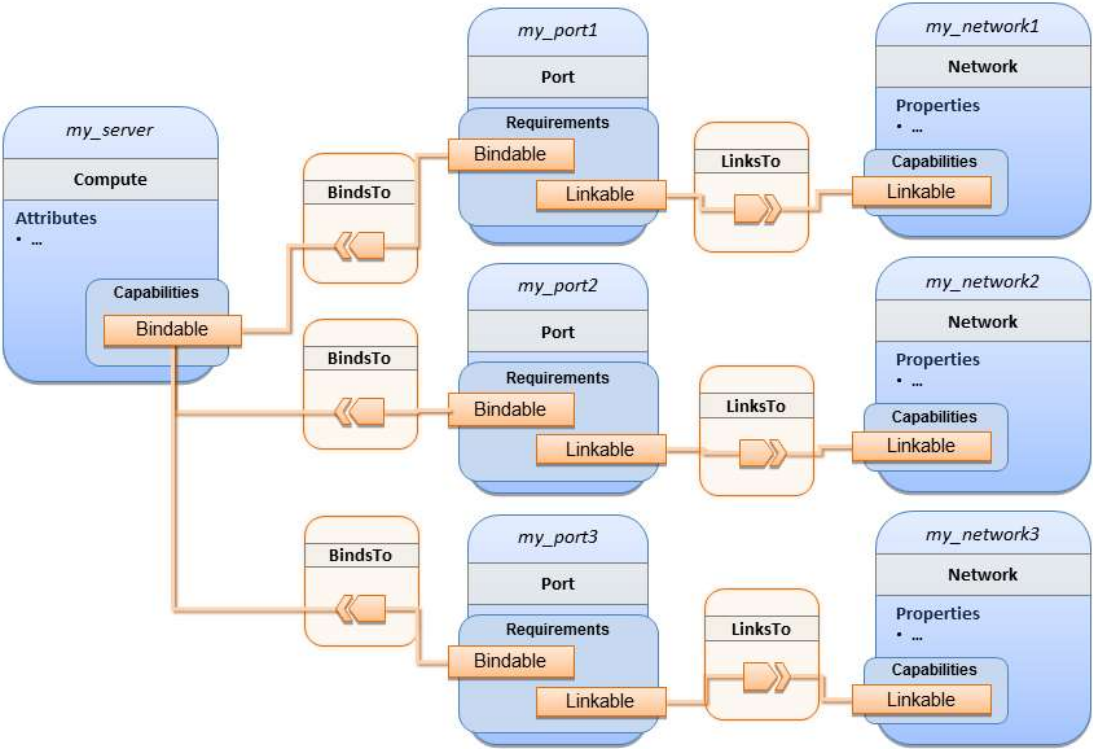
my_port2:
  type: toska.nodes.network.Port
  requirements:
    - binding: my_server2
    - link: my_network

```

3260 10.1.14 Network 4: Server bound to three networks

3261 10.1.14.1 Description

3262 This use case shows how three logical networks (Network), each with its own IP address range, can be bound to
 3263 with the same server (Compute node).



3265

3266 10.1.14.3 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: toska.nodes.network.Network
  properties:
    cidr: '192.168.1.0/24'
    network_name: net1

my_network2:
  type: toska.nodes.network.Network
  properties:
    cidr: '192.168.2.0/24'
    network_name: net2

my_network3:
  type: toska.nodes.network.Network
  properties:
    cidr: '192.168.3.0/24'
    network_name: net3

my_port1:
  type: toska.nodes.network.Port
  properties:
    order: 0
  requirements:
    - binding: my_server
    - link: my_network1

my_port2:
  type: toska.nodes.network.Port
  properties:
    order: 1
  requirements:
    - binding: my_server
    - link: my_network2

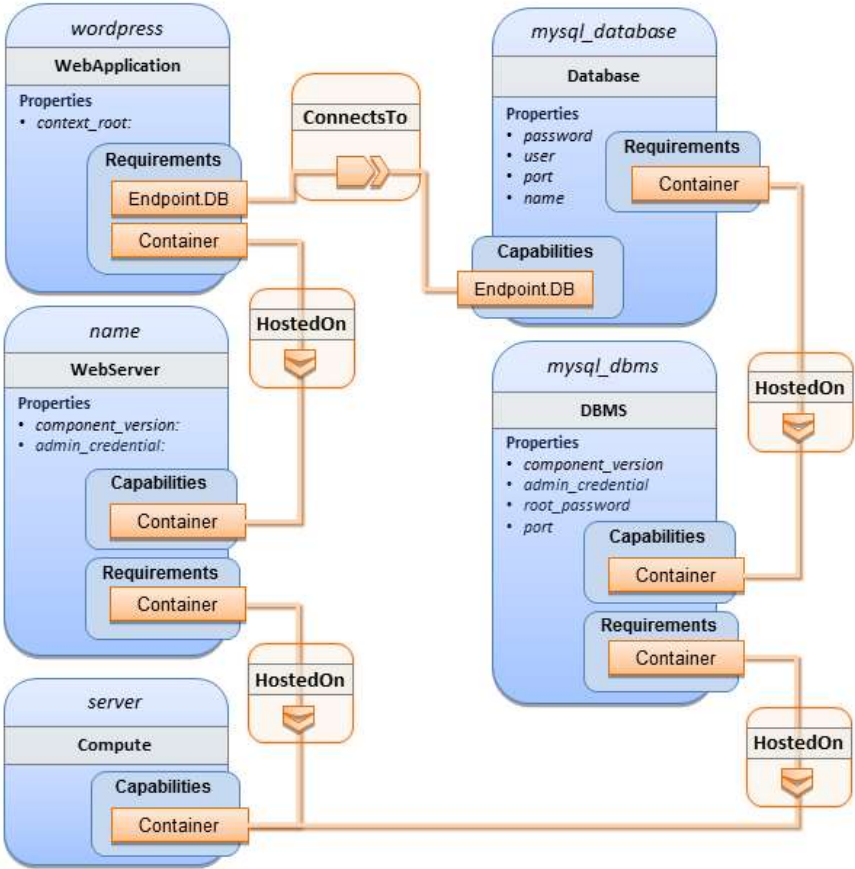
my_port3:
  type: toska.nodes.network.Port
  properties:
    order: 2
  requirements:
    - binding: my_server
    - link: my_network3
```

3267 **10.1.15 WebServer-DBMS 1: WordPress + MySQL, single instance**

3268 **10.1.15.1 Description**

3269 TOSCA simple profile service showing the WordPress web application with a MySQL database hosted on a single
3270 server (instance).

3271 **10.1.15.2 Logical Diagram**



3272

3273 **10.1.15.3 Sample YAML**

```
tosca_definitions_version: tosca_simple_yaml_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: tosa.nodes.WebApplication.WordPress
    properties:
      context_root: { get_input: context_root }
    requirements:
      - host: webserver
      - 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

webserver:
  type: WebServer
  requirements:
    - host: server
  interfaces:
    Standard:
      create: webserver\_install.sh
      start: webserver\_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] }
```

3274 **10.1.15.4 Sample scripts**

3275 Where the referenced implementation scripts in the example above would have the following contents

3276 **10.1.15.4.1 wordpress_install.sh**

```
yum -y install wordpress
```

3277 **10.1.15.4.2 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
```

3278 **10.1.15.4.3 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
```

3279 **10.1.15.4.4 mysql_dbms_install.sh**

```
yum -y install mysql mysql-server
# Use systemd to start MySQL server at system boot time
systemctl enable mysqld.service
```

3280 **10.1.15.4.5 mysql_dbms_start.sh**

```
# Start the MySQL service (NOTE: may already be started at image boot time)
systemctl start mysqld.service
```

3281 **10.1.15.4.6 mysql_dbms_configure**

```
# Set the MySQL server root password
mysqladmin -u root password db_root_password
```

3282 **10.1.15.4.7 webserver_install.sh**

```
yum -y install httpd
systemctl enable httpd.service
```

3283 **10.1.15.4.8 webserver_start.sh**

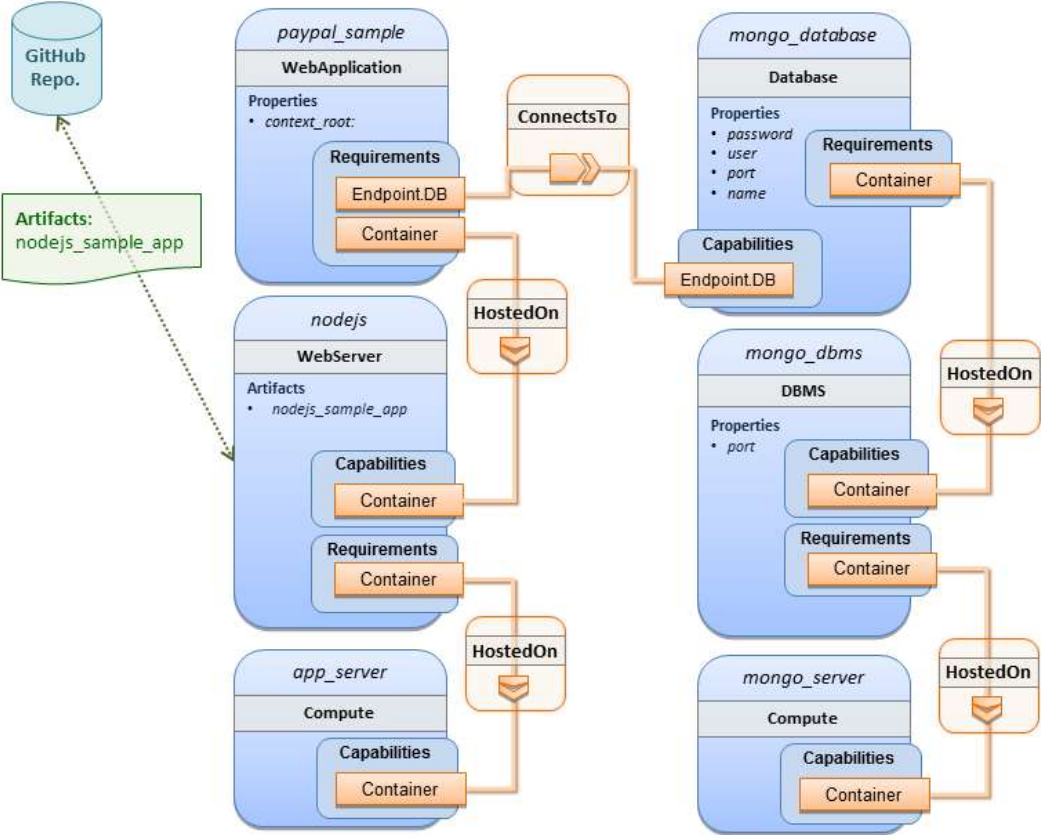
```
# Start the httpd service (NOTE: may already be started at image boot time)
systemctl start httpd.service
```

3284 **10.1.16 WebServer-DBMS 2: Nodejs with PayPal Sample App and MongoDB on**
3285 **separate instances**

3286 **10.1.16.1 Description**

3287 This use case Instantiates a 2-tier application with Nodejs and its (PayPal sample) WebApplication on one tier
3288 which connects a MongoDB database (which stores its application data) using a ConnectsTo relationship.

3289 10.1.16.2 Logical Diagram



3290

3291 10.1.16.3 Sample YAML

```
tosca_definitions_version: tosca_simple_yaml_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: scripts/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: toska.nodes.DBMS
  requirements:
    - host: mongo_server
  properties:
    port: 27017
  interfaces:
    toska.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: toska.nodes.Compute
  capabilities:
    os:
      properties: *os_capabilities
    host:
      properties: *ubuntu_node

app_server:
  type: toska.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 ] }

```

3292 **10.1.16.4 Notes:**

- 3293 • Scripts referenced in this example are assumed to be placed by the TOSCA orchestrator in the relative
3294 directory declared in TOSCA.meta of the TOSCA CSAR file.

3295 **10.1.17 Multi-Tier-1: Elasticsearch, Logstash, Kibana (ELK) use case with**
3296 **multiple instances**

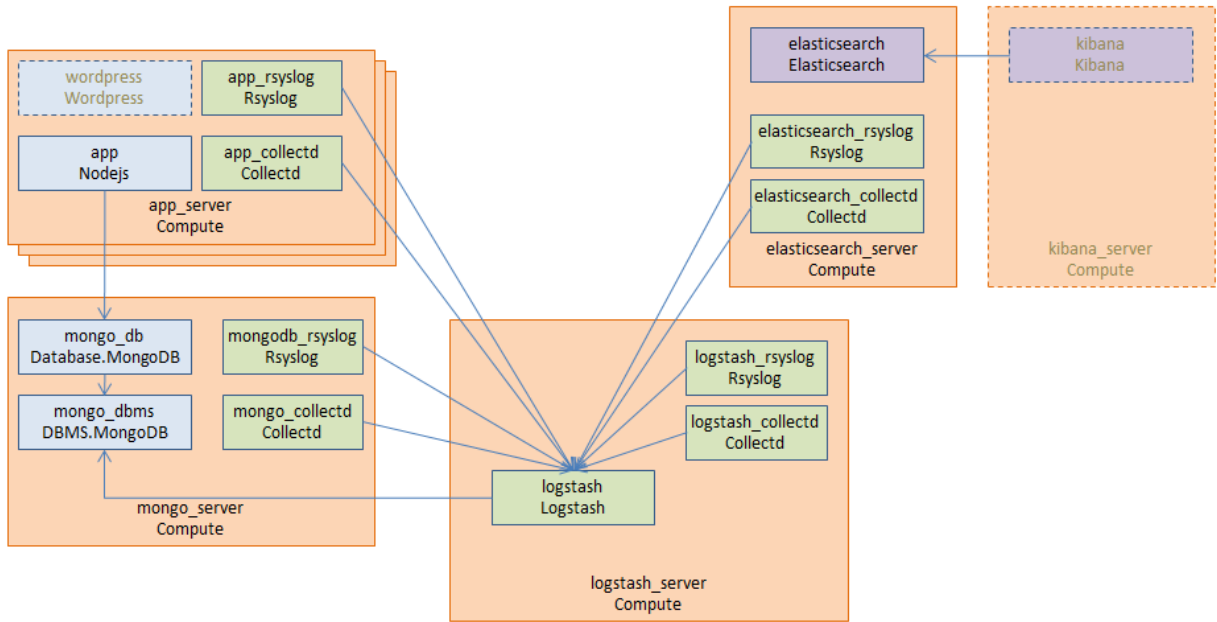
3297 **10.1.17.1 Description**

3298 TOSCA simple profile service showing the Nodejs, MongoDB, Elasticsearch, Logstash, Kibana, rsyslog and
3299 collectd installed on a different server (instance).

3300
3301 This use case also demonstrates:

- 3302 • Use of TOSCA macros or dsl_definitions
3303 • Multiple **SoftwareComponents** hosted on same Compute node
3304 • Multiple tiers communicating to each other over ConnectsTo using Configure interface.

3305 **10.1.17.2 Logical Diagram**



3306
3307 **10.1.17.3 Master Service Template application (Entry-Definitions)**

3308 TBD

3309

```
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: 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: Scripts/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
  interfaces:
    tosa.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: tosa.nodes.SoftwareComponent.Elasticsearch
  requirements:
    - host: elasticsearch_server
  interfaces:
    tosa.interfaces.node.lifecycle.Standard:
      create: Scripts/elasticsearch/create.sh

```

```

        start: Scripts/elasticsearch/start.sh
logstash:
  type: tosa.nodes.SoftwareComponent.Logstash
  requirements:
    - host: logstash_server
    - search_endpoint: elasticsearch
  interfaces:
    tosa.interfaces.relationship.Configure:
      pre_configure_source:
        implementation: Python/logstash/configure_elasticsearch.py
        input:
          elasticsearch_ip: { get_attribute: [elasticsearch_server,
ip_address] }
      interfaces:
        tosa.interfaces.node.lifecycle.Standard:
          create: Scripts/logstash/create.sh
          configure: Scripts/logstash/config.sh
          start: Scripts/logstash/start.sh

kibana:
  type: tosa.nodes.SoftwareComponent.Kibana
  requirements:
    - host: kibana_server
    - search_endpoint: elasticsearch
  interfaces:
    tosa.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: tosa.nodes.SoftwareComponent.Collectd
  requirements:
    - host: app_server
    - collectd_endpoint: logstash
  interfaces:
    tosa.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: tosca.nodes.Compute
  capabilities:
    host:
      properties: *host_capabilities
    os:
      properties: *os_capabilities

kibana_server:
  type: tosca.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 ] }

```

3310 10.1.17.4 Sample scripts

3311 Where the referenced implementation scripts in the example above would have the following contents

3312 **10.1.18 Container-1: Containers using Docker single Compute instance**
3313 **(Containers only)**

3314 **10.1.18.1 Description**

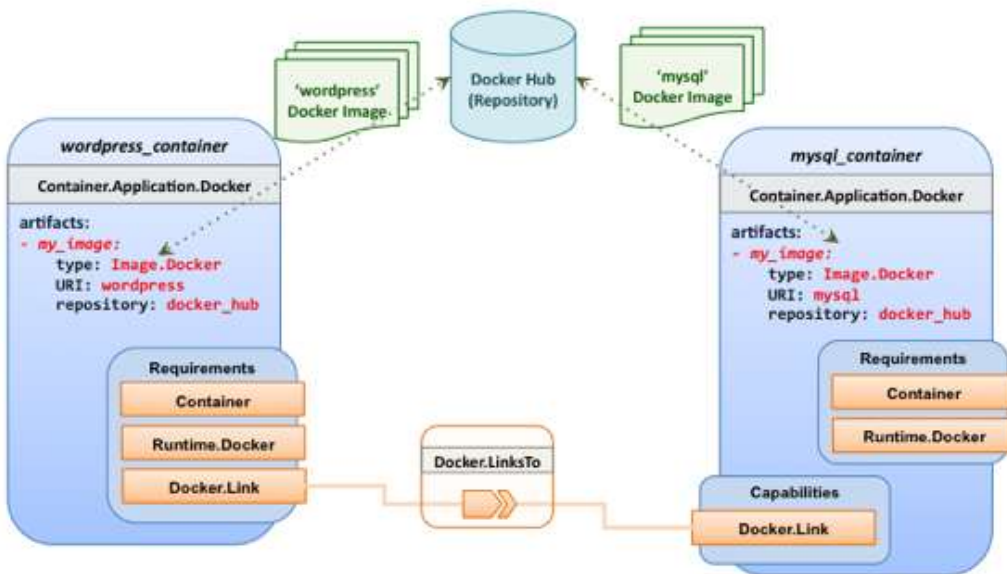
3315 This use case shows a minimal description of two Container nodes (only) providing their Docker Requirements
3316 allowing platform (orchestrator) to select/provide the underlying Docker implementation (Capability). Specifically,
3317 wordpress and mysql Docker images are referenced from Docker Hub.

3318

3319 This use case also demonstrates:

- 3320
- Abstract description of Requirements (i.e., Container and Docker) allowing platform to dynamically select the appropriate runtime Capabilities that match.
 - Use of external repository (Docker Hub) to reference image artifact.
- 3322

3323 **10.1.18.2 Logical Diagram**



3324

3325 **10.1.18.3 Sample YAML**

3326 **10.1.18.3.1 Two Docker “Container” nodes (Only) with Docker Requirements**

```
tosca_definitions_version: tosa_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 }

```

11 TOSCA Policies

This section is **non-normative** and describes the approach TOSCA Simple Profile plans to take for policy description with TOSCA Service Templates. In addition, it explores how existing TOSCA Policy Types and definitions might be applied in the future to express operational policy use cases.

11.1 A declarative approach

TOSCA Policies are a type of requirement that govern use or access to resources which can be expressed independently from specific applications (or their resources) and whose fulfillment is not discretely expressed in the application's topology (i.e., via TOSCA Capabilities).

TOSCA deems it not desirable for a declarative model to encourage external intervention for resolving policy issues (i.e., via imperative mechanisms external to the Cloud). Instead, the Cloud provider is deemed to be in the best position to detect when policy conditions are triggered, analyze the affected resources and enforce the policy against the allowable actions declared within the policy itself.

11.1.1 Declarative considerations

- Natural language rules are not realistic, too much to represent in our specification; however, regular expressions can be used that include simple operations and operands that include symbolic names for TOSCA metamodel entities, properties and attributes.
- Complex rules can actually be directed to an external policy engine (to check for violation) returns true/false then policy says what to do (trigger or action).
- Actions/Triggers could be:
 - Autonomic/Platform corrects against user-supplied criteria
 - External monitoring service could be utilized to monitor policy rules/conditions against metrics, the monitoring service could coordinate corrective actions with external services (perhaps Workflow engines that can analyze the application and interact with the TOSCA instance model).

11.2 Consideration of Event, Condition and Action

11.3 Types of policies

Policies typically address two major areas of concern for customer workloads:

- **Access Control** – assures user and service access to controlled resources are governed by rules which determine general access permission (i.e., allow or deny) and conditional access dependent on other considerations (e.g., organization role, time of day, geographic location, etc.).
- **Placement** – assures affinity (or anti-affinity) of deployed applications and their resources; that is, what is allowed to be placed where within a Cloud provider's infrastructure.
- **Quality-of-Service** (and continuity) - assures performance of software components (perhaps captured as quantifiable, measure components within an SLA) along with consideration for scaling and failover.

11.3.1 Access control policies

Although TOSCA Policy definitions could be used to express and convey access control policies, definitions of policies in this area are out of scope for this specification. At this time, TOSCA encourages organizations that already have standards that express policy for access control to provide their own guidance on how to use their standard with TOSCA.

3367 11.3.2 Placement policies

3368 There must be control mechanisms in place that can be part of these patterns that accept governance policies
3369 that allow control expressions of what is allowed when placing, scaling and managing the applications that are
3370 enforceable and verifiable in Cloud.

3371

3372 These policies need to consider the following:

- 3373 • Regulated industries need applications to control placement (deployment) of applications to different
3374 countries or regions (i.e., different logical geographical boundaries).

3375 11.3.2.1 Placement for governance concerns

3376 In general, companies and individuals have security concerns along with general “loss of control” issues when
3377 considering deploying and hosting their highly valued application and data to the Cloud. They want to control
3378 placement perhaps to ensure their applications are only placed in datacenter they trust or assure that their
3379 applications and data are not placed on shared resources (i.e., not co-tenanted).

3380

3381 In addition, companies that are related to highly regulated industries where compliance with government, industry
3382 and corporate policies is paramount. In these cases, having the ability to control placement of applications is an
3383 especially significant consideration and a prerequisite for automated orchestration.

3384 11.3.2.2 Placement for failover

3385 Companies realize that their day-to-day business must continue on through unforeseen disasters that might
3386 disable instances of the applications and data at or on specific data centers, networks or servers. They need to
3387 be able to convey placement policies for their software applications and data that mitigate risk of disaster by
3388 assuring these cloud assets are deployed strategically in different physical locations. Such policies need to
3389 consider placement across geographic locations as wide as countries, regions, datacenters, as well as granular
3390 placement on a network, server or device within the same physical datacenter. Cloud providers must be able to
3391 not only enforce these policies but provide robust and seamless failover such that a disaster’s impact is never
3392 perceived by the end user.

3393 11.3.3 Quality-of-Service (QoS) policies

3394 Quality-of-Service (apart from failover placement considerations) typically assures that software applications and
3395 data are available and performant to the end users. This is usually something that is measurable in terms of end-
3396 user responsiveness (or response time) and often qualified in SLAs established between the Cloud provider and
3397 customer. These QoS aspects can be taken from SLAs and legal agreements and further encoded as
3398 performance policies associated with the actual applications and data when they are deployed. It is assumed that
3399 Cloud provider is able to detect high utilization (or usage load) on these applications and data that deviate from
3400 these performance policies and is able to bring them back into compliance.

3401

3402 11.4 Policy relationship considerations

- 3403 • Performance policies can be related to scalability policies. Scalability policies tell the Cloud provider exactly
3404 **how** to scale applications and data when they detect an application’s performance policy is (or about to be)
3405 violated (or triggered).
- 3406 • Scalability policies in turn are related to placement policies which govern **where** the application and data can
3407 be scaled to.
- 3408 • There are general “tenant” considerations that restrict what resources are available to applications and data
3409 based upon the contract a customer has with the Cloud provider. This includes other constraints imposed by
3410 legal agreements or SLAs that are not encoded programmatically or associated directly with actual application
3411 or data..

3412 11.5 Use Cases

3413 This section includes some initial operation policy use cases that we wish to describe using the TOSCA
3414 metamodel. More policy work will be done in future versions of the TOSCA Simple Profile in YAML specification.

3415 11.5.1 Placement

3416 11.5.1.1 Use Case 1: Simple placement for failover

3417 11.5.1.1.1 Description

3418 This use case shows a failover policy to keep at least 3 copies running in separate containers. In this simple case,
3419 the specific containers to use (or name is not important; the Cloud provider must assure placement separation
3420 (anti-affinity) in three physically separate containers.

3421 11.5.1.1.2 Features

3422 This use case introduces the following policy features:

- 3423 • Simple separation on different “compute” nodes (up to discretion of provider).
- 3424 • Simple separation by region (a logical container type) using an allowed list of region names relative to the
3425 provider.
 - 3426 ○ Also, shows that set of allowed “regions” (containers) can be greater than the number of
3427 containers requested.

3428 11.5.1.1.3 Logical Diagram

3429 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
```

3430 11.5.1.1.4 Notes

- 3431 • There may be availability (constraints) considerations especially if these policies are applied to “clusters”.
- 3432 • There may be future considerations for controlling max # of instances per container.

3433 11.5.1.2 Use Case 2: Controlled placement by region

3434 11.5.1.2.1 Description

3435 This use case demonstrates the use of named “containers” which could represent the following:

- 3436 • Datacenter regions
- 3437 • Geographic regions (e.g., cities, municipalities, states, countries, etc.)
- 3438 • Commercial regions (e.g., North America, Eastern Europe, Asia Pacific, etc.)

3439 11.5.1.2.2 Features

3440 This use case introduces the following policy features:

- 3441 • Separation of resources (i.e., TOSCA nodes) by logical regions, or zones.

3442 11.5.1.2.3 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)
  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 ]
```

3443 11.5.1.3 Use Case 3: Co-locate based upon Compute affinity

3444 11.5.1.3.1 Description

3445 Nodes that need to be co-located to achieve optimal performance based upon access to similar Infrastructure
3446 (IaaS) resource types (i.e., Compute, Network and/or Storage).

3447
3448 This use case demonstrates the co-location based upon Compute resource affinity; however, the same approach
3449 could be taken for Network as or Storage affinity as well. :

3450 11.5.1.3.2 Features

3451 This use case introduces the following policy features:

- 3452 • Node placement based upon Compute resource affinity.

3453 11.5.1.4 Notes

- 3454 • The concept of placement based upon IaaS resource utilization is not future-thinking, as Cloud should
3455 guarantee equivalent performance of application performance regardless of placement. That is, all
3456 network access between application nodes and underlying Compute or Storage should have equivalent
3457 performance (e.g., network bandwidth, network or storage access time, CPU speed, etc.).

3458 11.5.1.4.1 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
```

3459 11.5.2 Scaling

3460 11.5.2.1 Use Case 1: Simple node autoscale

3461 11.5.2.1.1 Description

3462 Start with X nodes and scale up to Y nodes, capability to do this from a dashboard for example.

3463 **11.5.2.1.2 Notes**

3464 **11.5.2.1.3 Features**

3465 This use case introduces the following policy features:

- 3466
 - Basic autoscaling policy

3467 **11.5.2.1.4 Sample YAML**

```
my_scaling_policy_1:
  type: tosa.policy.scaling
  description: Simple node autoscaling
  properties:
    min_instances: <integer>
    max_instances: <integer>
    default_instances: <integer>
    increment: <integer>
```

3468 **11.5.2.1.5 Notes**

- 3469
 - Assume horizontal scaling for this use case
- 3470
 - Horizontal scaling, implies “stack-level” control using Compute nodes to define a “stack” (i.e., The
- 3471
 - Compute node's entire HostedOn relationship dependency graph is considered part of its “stack”)
- 3472
 - Assume Compute node has a SoftwareComponent that represents a VM application.
- 3473
 - Availability Zones (and Regions if not same) need to be considered in further use
- 3474
 - cases.
- 3475
 - If metrics are introduced, there is a control-loop (that monitors). Autoscaling is a special concept that
- 3476
 - includes these considerations.
- 3477
 - Mixed placement and scaling use cases need to be considered:
- 3478
 - Example: Compute1 and Compute2 are 2 node templates. Compute1 has 10 instances, 5 in one
- 3479
 - region 5 in other region.
- 3480

12 Conformance

12.1 Conformance Targets

The implementations subject to conformance are those intruded in Section 2.3 “Implementations”. They are listed here for convenience:

- TOSCA YAML service template
- TOSCA processor
- TOSCA orchestrator (also called orchestration engine)
- TOSCA generator
- TOSCA archive

12.2 Conformance Clause 1: TOSCA YAML service template

A document conforms to this specification as TOSCA YAML service template if it satisfies all the statements below:

- (a) It follows the grammar, rules and requirements defined in section 4 “TOSCA Simple Profile definitions in YAML”.
- (b) When using functions defined in section 5 “TOSCA functions”, it follows the grammar specified for these functions.
- (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 6 “TOSCA normative type definitions”, it follows the definitions given in this Section.

12.3 Conformance Clause 2: TOSCA processor

A processor or program conforms to this specification as TOSCA processor if it satisfies all the statements below:

- (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.
- (b) It implements the requirements and semantics associated with the definitions and grammar in section 4 “TOSCA Simple Profile definitions in YAML”, including those listed in the “additional requirements” subsections.
- (c) It is capable of resolving the imports, either explicit or implicit, as described in section 4 “TOSCA Simple Profile definitions in YAML”.

12.4 Conformance Clause 3: TOSCA orchestrator

A processor or program conforms to this specification as TOSCA orchestrator if it satisfies all the statements below:

- (a) It is conforming as a TOSCA Processor.
- (b) It is capable of using artifacts of the types described in section 6.3 “Artifact types”.
- (c) It is capable of processing TOSCA archives as intended in section 7 “TOSCA Cloud Service Archive (CSAR) format” and other related normative sections.
- (d) It is capable of understanding and processing the functions defined in section 5 “TOSCA functions” according to their rules and semantics.
- (e) It is capable of understanding and processing the normative type definitions according to their semantics and requirements as described in section 6 “TOSCA normative type definitions”.
- (f) It is capable of understanding and processing the networking types and semantics defined in section 8 “TOSCA Networking”.

3523

3524 **12.5 Conformance Clause 4: TOSCA generator**

3525 A processor or program conforms to this specification as TOSCA generator if it satisfies all the statements below:

- 3526 (a) It is capable of producing a conforming TOSCA YAML service template.
3527 (b) It does not produce documents that do not conform as TOSCA YAML service template while clearly
3528 intending to.

3529

3530 **12.6 Conformance Clause 5: TOSCA archive**

3531 A package artifact conforms to this specification as TOSCA archive if it satisfies all the statements below:

- 3532 (a) It follows the structure and rules defined in section 7 “TOSCA Cloud Service Archive (CSAR) format”.
3533

Appendix A. Known Extensions to TOSCA v1.0

The following items will need to be reflected in the TOSCA (XML) specification to allow for isomorphic mapping between the XML and YAML service templates.

A.1 Model Changes

- The “TOSCA Simple ‘Hello World’” example introduces this concept in Section 2. Specifically, a VM image assumed to be accessible by the cloud provider.
- Introduce template Input and Output parameters
- The “Template with input and output parameter” example introduces concept in Section 2.1.1.
 - “Inputs” could be mapped to BoundaryDefinitions in TOSCA v1.0. Maybe needs some usability enhancement and better description.
 - “outputs” are a new feature.
- Grouping of Node Templates
 - This was part of original TOSCA proposal, but removed early on from v1.0. This allows grouping of node templates that have some type of logically managed together as a group (perhaps to apply a scaling or placement policy).
- Lifecycle Operation definition independent/separate from Node Types or Relationship types (allows reuse). For now we added definitions for “node.lifecycle” and “relationship.lifecycle”.
- Override of Interfaces (operations) in the Node Template.
- Service Template Naming/Versioning
 - Should include TOSCA spec. (or profile) version number (as part of namespace)
- Allow the referencing artifacts using a URL (e.g., as a property value).
- Repository definitions in Service Template.
- Substitution mappings for Topology template.
- Addition of Group Type, Policy Type, Group def., Policy def. along with normative TOSCA base types for policies and groups.

A.2 Normative Types

- Constraints
 - constraint clauses, regex
- Types / Property / Parameters
 - list, map, range, scalar-unit types
 - Includes YAML intrinsic types
 - NetworkInfo, PortInfo, PortDef, PortSpec, Credential
 - TOSCA Version based on Maven
- Node
 - Root, Compute, ObjectStorage, BlockStorage, Network, Port, SoftwareComponent, WebServer, WebApplication, DBMS, Database, Container, and others
- Relationship
 - Root, DependsOn, HostedOn, ConnectsTo, AttachesTo, RoutesTo, BindsTo, LinksTo and others
- Artifact
 - Deployment: Image Types (e.g., VM, Container), ZIP, TAR, etc.
 - Implementation: File, Bash, Python, etc.
- Requirements
 - None
- Capabilities
 - Container, Endpoint, Attachment, Scalable, ...

3579 • Lifecycle
3580 • Standard (for Node Types)
3581 • Configure (for Relationship Types)
3582 • Functions
3583 • get_input, get_attribute, get_property, get_nodes_of_type, get_operation_output and others
3584 • concat, token
3585 • get_artifact
3586 • Groups
3587 • Root
3588 • Policies
3589 • Root, Placement, Scaling, Update, Performance
3590

Appendix B. Acknowledgments

The following individuals have participated in the creation of this specification and are gratefully acknowledged:

Contributors:

Avi Vachnis (avi.vachnis@alcatel-lucent.com), Alcatel-Lucent
Chris Lauwers (lauwers@ubicity.com)
Derek Palma (dpalma@vnomic.com), Vnomic
Frank Leymann (Frank.Leymann@informatik.uni-stuttgart.de), Univ. of Stuttgart
Gerd Breiter (gbreiter@de.ibm.com), IBM
Hemal Surti (hsurti@cisco.com), Cisco
Ifat Afek (ifat.afek@alcatel-lucent.com), Alcatel-Lucent
Idan Moyal, (idan@gigaspace.com), Gigaspaces
Jacques Durand (jdurand@us.fujitsu.com), Fujitsu
Jin Qin, (chin.qinjin@huawei.com), Huawei
Juergen Meynert (juergen.meynert@ts.fujitsu.com), Fujitsu
Kapil Thangavelu (kapil.thangavelu@canonical.com), Canonical
Karsten Beins (karsten.beins@ts.fujitsu.com), Fujitsu
Kevin Wilson (kevin.l.wilson@hp.com), HP
Krishna Raman (kraman@redhat.com), Red Hat
Luc Boutier (luc.boutier@fastconnect.fr), FastConnect
Matt Rutkowski (mrutkows@us.ibm.com), IBM
Moshe Elisha (moshe.elisha@alcatel-lucent.com), Alcatel-Lucent
Nate Finch (nate.finch@canonical.com), Canonical
Nikunj Nemani (nnemani@vmware.com), VMware
Richard Probst (richard.probst@sap.com), SAP AG
Sahdev Zala (spzala@us.ibm.com), IBM
Shitao Li (lishitao@huawei.com), Huawei
Simeon Monov (sdmonov@us.ibm.com), IBM
Stephane Maes (stephane.maes@hp.com), HP
Thomas Spatzier (thomas.spatzier@de.ibm.com), IBM
Ton Ngo (ton@us.ibm.com), IBM
Travis Tripp (travis.tripp@hp.com), HP
Vahid Hashemian (vahidhashemian@us.ibm.com), IBM
Wayne Witzel (wayne.witzel@canonical.com), Canonical
Yaron Parasol (yaropa@gigaspace.com), Gigaspaces

Appendix C. Revision History

Revision	Date	Editor	Changes Made
WD06, Rev01	2015-06-02	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • Initial WD06, Revision 01 baseline. • Ch. 3: Changed Hello World example to use 512 MB instead of 4 MB which was not a realistic value. • C.8.4: Fixed incorrect Heading number (indent) • C.8.1: added “version” property to <code>tosca.artifacts.Root</code> to allow template authors the ability to declare the artifact’s version (i.e., the file itself and its contents). • H.1.2: Adjusted description of Compute 1 use case to be more accurate • H.1.3: Added new SoftwareComponent use case to show how an existing VM image can be used as a deployment artifact. • H.1.19: Container 1: Fixed incorrect YAML which did not use the artifact definition grammar correctly • A.5.5: Artifact def.: Added missing “repository” keyname to grammar and renamed “implementation” keyname to “file” to simplify and to not confuse template authors since this keyname is also used for deployment artifacts. • C.2.1: Added <code>tosca.datatypes.Root</code> that all TOSCA complex datatypes derive from • C.2.2, C.2.3, C.2.4, C.2.6: Credential, NetworkInfo, PortInfo, PortSpec now derive from <code>tosca.datatypes.Root</code>. • A.6.3, A.6.4, A.6.5, A.6.6, A.6.8, A.6.9: Added optional “version” keyname to all TOSCA type grammars (i.e., Artifact, Interface, Data, Capability, Node and Relationship). • A.6.4: Interface Type: Added optional “description” keyname. • A.1.1, A.1.2: Added Namespace prefix, Namespacing in TOSCA Service Templates sections. • A.1.1.3: Updated namespace requirements/rules. • A.9.1: Removed keyname “<code>tosca_default_namespace</code>” since it was redundant to the one indicated by the “<code>tosca_definitions_version</code>” keyname. • A.9.3.1: Fixed namespace errors in Service Template examples. • A.7.4: Removed “alias” keyname from Relationship Template as we had no use cases for it. Also, fixed copy to be a “string” type. • A.5.7: Added “import definition” section with keynames, grammars and example. • A.9: Service Template: Fixed “imports” grammar and example.
WD06, Rev02	2015-06-02	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • D.1.3: Added QCOW2 as non-normative VM image format • D.1.2, D.1.3: Added “mime_type” and “file_ext” values to the ISO and QCOW2 non-normative artifact types. • H.1.3: Changed use case example to use QCOW2 (which is testable on various hypervisors such as XEN and KVM). • H.1.3: Added assumptions to the use case to assure that this is qualified as a non-baremetal use case and that VM management services (including hypervisors) are in place that support the VM image format. • Fixed all hyperlinks to any <code>tosca.capabilities.Endpoint.*</code> capability definitions. • C.8.3, C.8.4: Added “Additional Requirements” to allow orchestrators to throw an error if they encounter non-normative deployment or implementation artifacts they cannot process.
WD06, Rev03	2015-06-22	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • A.5.7.4: import def.: fixed example to have real sample values for repository, namespace URI and namespace prefix. • A.6.10: Added Group Type definition, grammar and example • A.7.5: Added Group Definition and example • 14: Fixed example in group chapter to use latest grammar. • A.7.4: Added missing Relationship Def. example • C.8: Artifact Type: Fixed description to be more accurate • 11.3: Included example of how to reference an abstract database without

			<p>a node filter (to show alternative using an abstract node template in the topology).</p> <ul style="list-style-type: none"> • A.5: removed redundant word “recognized” from all keyname table introductions. • A.8: Adjusted topology definition to reference new Group definition grammar. • A.9: Service Template: Added support for “group_types” element and listing Group Type definitions. • Appendix I: Policies: Added official prologue on TOSCA’s declarative approach stated the areas of policy we acknowledge include Access Control, Governance (placement) and Quality-of-Service (QoS).
WD06, Rev04	2015-07-01	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • A.6.10: Group Type: Added up-front discussion of the distinction between a group and a node In TOSCA. • A.6.10, A.7.5: Group Type, Group definition: Added “interfaces” • C.9: Developed the normative Group Type section more. • C.9: Added Standard interface to toscagroups.Root Type. • I3: Lots of prose to discuss types of policies and considerations (and overlap). • I.5: Policy Use Cases: Added use cases section and started to organize our policy use cases by policy type (e.g., placement, QoS, scaling, etc.) • I.5.1, I.5.2: Worked on separating out individual Placement and scaling use cases for policies.
WD06, Rev05	2015-07-08	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • 11.2: Simple Profile WG agreed to add back the “node_filter” keyword and ordered lists for capabilities for “abstract” nodes; otherwise, ordering was lost for requirement fulfillment. Additionally, we agreed that the grammar needs to support “constraints” wherever they appear on abstract nodes. • A.5.4.5: Removed the portion of the example that hinted at the use of a capability “decorator” in the form of requesting encryption in the CPU. We kept the removed capability as a comment to address in v1.1. • D.3.1.2: Database.MySQL: does not need “root_password” property. This was likely a copy/paste error from DBMS type that no one caught until now. • A.5.5.1 Artifact def. – prose in table still used old “implementation” keyname; change it to current keyname of “file”. • Ch 6: Changed artifact definition example to use “file” keyname to match grammar. • A.5.12: Added comment to Operation Def. to discuss using “file” keyname instead of “implementation” which would be simpler. • A.5.12.2.3: Fixed Operation Def. grammar when multiple files are needed to implement the operation (i.e., “implementation” keyname has its own map of keynames). • 11.1: Missing colon after “host” capability • 11.3: Missing colon after “my_abstract_database” • 14: Removed “policies” from “group” example as this is not specified and will not be how we associate policies with groups. • A.2.4.2: typo on constraint in example. • A.5.5: Artifact def. – “file” is now a required keyname, resolves issue TOSCA-249. • A.5.5.2: Fixed typos for repository definition and fixed font as well. • A.5.12.4.1: Fixed single-line example for start operation def. • A.5.12.4.2: Fixed example so that secondary artifacts are in a valid YAML ordered list format. • A.6.3: Fixed description for Artifact Type to be coherent. • A.6.8.2: Fixed copy/paste error for artifact def. description below grammar. • A.6.10, A.7.5: Added “properties” keyname to Group Type and Group definition along with their respective property definitions and assignments. • A.8: Topology Template: formalized grammar so it is like all other definitions. Clarified use of groups within topology. Comments to create a new parameter definition for inputs and outputs. • A.9.3: Added subsections under Service Template to describe “group_types”, “artifact_types”, “relationships” and “interface_types”.

			<ul style="list-style-type: none"> • A.9.10: Renamed “datatype_definitions” to “data_types” • A.9.1: Changed Service Template grammar to use the new Import Definition (no longer only a list of simple strings). • C.2.2: Credential: Table had listed the “protocol” as required property which it is not. • C.2.2: Credential: Fixed examples to remove “type” and “properties” keys that are not valid for a property assignment (only property def.). • C.2.5.2: Missing colon from example for “in_range” constraint • C.2.6: PortSpec: “target” and “source” grammars are not maps but a single PortDef. • C.8.1: tosca.artifacts.Root: Does not need a version property as it is now part of the metamodel. • C.9.1: tosca.groups.Root: Does not need a version property as it is now part of the metamodel. • D.1.3: QCOW2 artifact type had “ISO” in type name; a cut/paste error. • H.1.3.5: Use case YAML had an artifact as an ordered list; and “host” requirement was not ordered list. Fixed both. • G.1.4, H.1.4, H.1.5, H.1.7, H.1.8, and H.1.9: BlockStorage use cases should use “local_storage” requirement already pre-defined in the Compute node not “attachment” which is not normative. • H.1.5, H.1.6: Compute node needs to reference “my_storage” node. • H.1.6: Indentation error on one of the output parameters. • H.1.9: Incorrect YAML on “local_storage” relationship, AttachesTo type needs to be indented on a new line with “type” from Extended Grammar as explained in section A.7.2. • H.1.12: Added “host” capability and moved its properties that were directly under Compute node (old grammar) to be under “host”. Also added units to property values. • H.1.19: “inputs” definitions were not inside topology template (old grammar). Fixed that the image artifacts were defined as ordered lists, normal unordered list now. Also, only need interface name “Standard” not redeclared the normative type definition.
WD06, Rev06	2015-07-28	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • Changed all examples to use a host container mem_size of either 512 MB (i.e., “tiny”) or 4096 MB (i.e., medium) to be more realistic to typical production public cloud values. • A.5.12: Added Parameter def. which is used by Input and Output parameters for the topology template and Inputs for Interfaces • A.6.10: Added Policy Type section, grammar and examples • A.7.6: Added Policy Def. section, grammar and examples • C.10 Policy Types: Added normative policy type (families) which will be expanded in future versions. The goal in v1.0 was to establish the type families acknowledged in Appendix I. These include Root, Placement, Scaling, Performance and Update. • C.6.6: RoutesTo relationship type: Removed empty grammar and example sections. • A.8.2: Topology Template: Fixed grammar issues with group definitions. Added policy definitions to grammar. Changed grammar to reference Parameter def. for both Inputs and Outputs. • A.8.2.1: Topology Template: Changed Inputs section to reference Parameter def. • A.8.2.4: Topology Template: Changed Outputs section to reference Parameter def. • A.8.2.6: Topology Template: Added policies section. • A.9.1: Service Template: Added “policies” keyname section to allow policy types to be declared. • H.1.15.4.3: Changed legacy “postconfigure” operations to current “configure” operation.
WD06, Rev07	2015-08-20	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • A.9.1: Service Template: Added metadata keyname “mime_type” to declare the type of TOSCA file’s Mime Type string (i.e. grammar, information) the .yaml file holds. • A.6.10, A.6.11: Added the keyname “targets” to both Group Type and Policy Type as a new “shortcut” for describing a relationship between a logical entity (i.e. a Group or Policy) and concrete entities in TOSCA models (i.e., ones that are part of the actual application topology and

			<p>deployment such as Nodes).</p> <ul style="list-style-type: none"> • I.3.2.1: Placement policy: corrected incomplete sentence. • C.6.1: Relationship Types: added “state” attribute to Root type • A.3.2, C.6.1: Relationship States: Placeholder created to define states once TOSCA-272 is addressed. • I.5.1.3: Filled out “co-locate” placement policy use case removing TBD placeholders. • D.3: Non-normative Node Types: Added text to clearly state these are non-normative for use in examples and use cases and any Attributes, properties, requirements or capabilities they introduce are also non-normative. Fixed Properties tables to show “N/A” and opened TOSCA-273 if in the future we wish to document non-normative properties. • H.1, H.1.16: Removed the placeholder for the WebServer-DBMS 2 use case (will renumber use case WebServer-DBMS 3 to be 2) As it was a legacy AWS Cloud formation (CFN) use case that used explicit Floating IPs which is not the approach we use in TOSCA.
WD06, Rev08	2015-08-26	Matt Rutkowski, IBM Jacques Durand, Fujitsu	<ul style="list-style-type: none"> • Reorganization of chapters and appendices to conform to OASIS staff directions and to accommodate addition of a new Conformance section. <ul style="list-style-type: none"> ○ Created new Chapter 1 “Introduction”. <ul style="list-style-type: none"> ▪ Moved former Chapter 1 “Objective” to Chapter 1.1. ▪ Moved former Chapter 2 “Summary of key TOSCA Concepts” to Chapter 1.2. ▪ Created Chapter 1.3 “Implementations” to introduce key terms used to identify conformance targets for the new Conformance Chapter (12). ▪ Moved former Appendix J.2 “Terminology” to Chapter 1.4. ▪ Moved RFC 2119 text formerly under Appendix J.2 under a new Chapter 1.5 “Notational Conventions” adding a normative reference to RFC 2119. ▪ Moved former Appendix J.3, J.4 “Normative References” and “Non-normative References” to Chapters 1.6 and 1.7 respectively. ▪ Moved former Appendix J.5 “Glossary” to Chapter 1.8. ○ Created new Chapter 2 “TOSCA by example”. <ul style="list-style-type: none"> ▪ Moved former Chapters 3 through 18 under new Chapter 2. This effectively made these chapters subchapters/subsections under the new chapter (which demotes all former Header levels by one decimal point). For example former Chapter 3 ‘A “hello world” template for TOSCA’ became Chapter 2.1 ‘A “hellow world” template for TOSCA’ and so on. ○ Renumbered former Appendix A “TOSCA Simple Profile definitions in YAML” to now be Chapter 3. <ul style="list-style-type: none"> ▪ Note: This chapter begins the group of normative chapters. ○ Renumbered former Appendix B “Functions” to now be Chapter 4. <ul style="list-style-type: none"> ▪ Changed title to “TOSCA functions” for consistency. ○ Renumbered former Appendix C “TOSCA normative type defintions” to now be Chapter 5. ○ Renumbered former Appendix D “Non-normative type defintions” to now be Chapter 8. <ul style="list-style-type: none"> ▪ Note: This new Chapter 8 locates (groups) this section with the other sections that are non-normative. This was done to accommodate the new “Conformance” section (Chapter 12) and its referential contents. ○ Renumbered former Appendix E “TOSCA Cloud Service Archive (CSAR)” to now be Chapter 6. ○ Renumbered former Appendix F “Networking” to now be Chapter 7. <ul style="list-style-type: none"> ▪ Changed title to “TOSCA networking” for consistency. ○ Renumbered former Appendix G “Component modeling use cases” to now be Chapter 9. <ul style="list-style-type: none"> ▪ Note: This chapter begins the group of non-normative chapters. ○ Renumbered former Appendix H “Complete application modeling use cases” to now be Chapter 10. <ul style="list-style-type: none"> ▪ Changed title to “Application modeling use cases” for consistency. ○ Renumbered former Appendix I “TOSCA Policies” to now be Chapter 11. ○ Moved contents of former Appendix J.1 “Known extensions to TOSCA

			<p>v1.0" to now be Appendix A contents.</p> <ul style="list-style-type: none"> ○ Former Appendix K "Issues list" was removed by consensus as it was outdated and should not appear in the final CSD 04 contents and be carried as the spec. is taken through the OASIS standards track process. <ul style="list-style-type: none"> ▪ Note: The "Issues List" will be added back for the first working draft of v1.1 of Simple Profile and updated to reflect the TOSCA JIRA status of issues against this work product. ○ Former Appendix L "Acknowledgments" was automatically renumbered to be Appendix B. ○ Former Appendix M "Revision history" was automatically renumbered to be Appendix C. • Chapter 1.4: Jacques authored "Implementations" section to reflect terms we use later in the Conformance section • Chapter 1.5: Authored "Terminology" introductory paragraph to explain the domain of TOSCA. • Chapter 12: "Conformance": Jacques authored this section which is a new OASIS requirement for standards-track documents. It establishes the Simple Profile in YAML conformance clauses relative to the Terms established in Section 1.5 and also by referencing the Chapters' contents. • Added a sentence at the top (i.e., first sentence) of every Chapter stating whether to Chapter contents was "normative" or "non-normative". • Section 10.1.18: Container-1 example: Fixed requiriements and capability for "linking" Docker containers. <ul style="list-style-type: none"> ○ Added logical diagram. • Section 3: Assured all Examples had captions (and that they appear in Table of Figures). • Section 4.4.2.3: Added get_property with TARGET parameter. • Section 2.10.3: Changed 'SHOULD' to 'should'. • Section 3.8.2: Topology template: Policies should be a sequenced list. • Section 3.5.13.2: Operation def.: Fixed extened grammar which was missing 'primary' keyword. • Section 3.6.10.3: Group Type: Fixed example to use 'targets'. • Section 3.6.11: Polcy Type: Updated description for 'target' keyname to include Group types. • Section 3.7.5: Group def.: Changed 'members' to 'targets' as agreed by work group. • Section 2.12: Example 21: Fixed group example to adhere to latest grammar. • Section 3.7.6: Policy def.: Added 'targets' keyname as agreed by work group. • Section 3.8.1: Topology Template: Fixed description for 'inputs' and 'output' to reference parameter def. instead of property def. • Section 3.8.2.5.2: Topology Template: Fixed example to adhere to latest Group def. grammar. • Section 3.9.3.17: Added missing Policy Type defintions from Service Template defintion grammar. • Fixed hyperlinks throughout document • Chapter H.1: Filled in missing links to BlockStorage use cases. Fixed Table width to be 100% not 115%. • Incorporated several copy/paste errors, typos and such reported by Luc B.
--	--	--	---