

# TOSCA Simple Profile in YAML Version 1.1

## Committee Specification Draft 01 / Public Review Draft 01

25 August 2016

### Specification URIs

#### This version:

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

#### Previous version:

N/A

#### Latest version:

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

#### Technical Committee:

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

#### Chairs:

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

#### Editors:

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

#### Related work:

This specification is related to:

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

#### Declared XML namespace:

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

#### Abstract:

This document defines a simplified profile of the TOSCA Version 1.0 specification in a YAML rendering which is intended to simplify the authoring of TOSCA service templates. This profile defines a less verbose and more human-readable YAML rendering, reduced level of indirection between different modeling artifacts as well as the assumption of a base type system.

**Status:**

This document was last revised or approved by the 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](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-v1.1]**

*TOSCA Simple Profile in YAML Version 1.1*. Edited by Matt Rutkowski and Luc Boutier. 25 August 2016. OASIS Committee Specification Draft 01 / Public Review Draft 01. <http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.1/csprd01/TOSCA-Simple-Profile-YAML-v1.1-csprd01.html>. Latest version: <http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.1/TOSCA-Simple-Profile-YAML-v1.1.html>.

---

## Notices

Copyright © OASIS Open 2016. All Rights Reserved.

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

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

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

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

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

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

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

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

---

# Table of Contents

Table of Examples .....	7
Table of Figures .....	7
1 Introduction .....	8
1.1 Objective .....	8
1.2 Summary of key TOSCA concepts .....	8
1.3 Implementations .....	8
1.4 Terminology .....	9
1.5 Notational Conventions .....	9
1.6 Normative References .....	9
1.7 Non-Normative References .....	10
1.8 Glossary .....	10
2 TOSCA by example .....	12
2.1 A “hello world” template for TOSCA Simple Profile in YAML .....	12
2.2 TOSCA template for a simple software installation .....	14
2.3 Overriding behavior of predefined node types .....	16
2.4 TOSCA template for database content deployment .....	17
2.5 TOSCA template for a two-tier application .....	19
2.6 Using a custom script to establish a relationship in a template .....	22
2.7 Using custom relationship types in a TOSCA template .....	23
2.8 Defining generic dependencies between nodes in a template .....	25
2.9 Describing abstract requirements for nodes and capabilities in a TOSCA template .....	26
2.10 Using node template substitution for model composition .....	30
2.11 Using node template substitution for chaining subsystems .....	34
2.12 Grouping node templates .....	40
2.13 Using YAML Macros to simplify templates .....	43
2.14 Passing information as inputs to Nodes and Relationships .....	44
2.15 Topology Template Model versus Instance Model .....	45
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
3.2 Parameter and property types .....	49
3.3 Normative values .....	58
3.4 TOSCA Metamodel .....	60
3.5 Reusable modeling definitions .....	60
3.6 Type-specific definitions .....	89
3.7 Template-specific definitions .....	107
3.8 Topology Template definition .....	119
3.9 Service Template definition .....	125
4 TOSCA functions .....	138
4.1 Reserved Function Keywords .....	138
4.2 Environment Variable Conventions .....	138
4.3 Intrinsic functions .....	141
4.4 Property functions .....	142

4.5	Attribute functions .....	145
4.6	Operation functions.....	146
4.7	Navigation functions .....	147
4.8	Artifact functions .....	147
4.9	Context-based Entity names (global) .....	150
5	TOSCA normative type definitions .....	151
5.1	Assumptions .....	151
5.2	TOSCA normative type names.....	151
5.3	Data Types.....	151
5.4	Artifact Types.....	160
5.5	Capabilities Types .....	163
5.6	Requirement Types .....	174
5.7	Relationship Types .....	174
5.8	Interface Types .....	177
5.9	Node Types.....	183
5.10	Group Types .....	196
5.11	Policy Types .....	196
6	TOSCA Cloud Service Archive (CSAR) format.....	198
6.1	Overall Structure of a CSAR.....	198
6.2	TOSCA Meta File.....	198
6.3	Archive without TOSCA-Metadata.....	199
7	TOSCA workflows .....	200
7.1	Normative workflows.....	200
7.2	Declarative workflows .....	200
7.3	Imperative workflows .....	204
7.4	Making declarative more flexible and imperative more generic .....	219
8	TOSCA networking.....	222
8.1	Networking and Service Template Portability.....	222
8.2	Connectivity Semantics .....	222
8.3	Expressing connectivity semantics.....	223
8.4	Network provisioning .....	225
8.5	Network Types.....	229
8.6	Network modeling approaches .....	235
9	Non-normative type definitions.....	240
9.1	Artifact Types.....	240
9.2	Capability Types .....	240
9.3	Node Types.....	242
10	Component Modeling Use Cases.....	246
11	Application Modeling Use Cases.....	253
11.1	Use cases .....	253
12	TOSCA Policies.....	301
12.1	A declarative approach.....	301
12.2	Consideration of Event, Condition and Action .....	301
12.3	Types of policies .....	301
12.4	Policy relationship considerations .....	302

12.5 Use Cases .....	303
13 Conformance .....	306
13.1 Conformance Targets .....	306
13.2 Conformance Clause 1: TOSCA YAML service template .....	306
13.3 Conformance Clause 2: TOSCA processor .....	306
13.4 Conformance Clause 3: TOSCA orchestrator .....	306
13.5 Conformance Clause 4: TOSCA generator .....	307
13.6 Conformance Clause 5: TOSCA archive .....	307
Appendix A. Known Extensions to TOSCA v1.0 .....	308
A.1 Model Changes .....	308
A.2 Normative Types .....	308
Appendix B. Acknowledgments .....	310
Appendix C. Revision History .....	311

---

## Table of Examples

Example 1 - TOSCA Simple "Hello World" .....	12
Example 2 - Template with input and output parameter sections .....	13
Example 3 - Simple (MySQL) software installation on a TOSCA Compute node .....	15
Example 4 - Node Template overriding its Node Type's "configure" interface .....	16
Example 5 - Template for deploying database content on-top of MySQL DBMS middleware .....	17
Example 6 - Basic two-tier application (web application and database server tiers).....	20
Example 7 - Providing a custom relationship script to establish a connection .....	22
Example 8 - A web application Node Template requiring a custom database connection type .....	24
Example 9 - Defining a custom relationship type.....	25
Example 10 - Simple dependency relationship between two nodes.....	25
Example 11 - An abstract "host" requirement using a node filter .....	27
Example 12 - An abstract Compute node template with a node filter.....	28
Example 13 - An abstract database requirement using a node filter .....	29
Example 14 - An abstract database node template .....	30
Example 15 - Referencing an abstract database node template.....	32
Example 16 - Using substitution mappings to export a database implementation .....	33
Example 17 - Declaring a transaction subsystem as a chain of substitutable node templates .....	35
Example 18 - Defining a TransactionSubsystem node type .....	37
Example 19 - Implementation of a TransactionSubsystem node type using substitution mappings.....	39
Example 20 - Grouping Node Templates for possible policy application .....	40
Example 21 - Grouping nodes for anti-colocation policy application .....	41
Example 22 - Using YAML anchors in TOSCA templates .....	43
Example 23 - Properties reflected as attributes .....	46

---

## Table of Figures

Figure 1: Using template substitution to implement a database tier .....	31
Figure 2: Substitution mappings .....	33
Figure 3: Chaining of subsystems in a service template .....	35
Figure 4: Defining subsystem details in a service template.....	38
Figure-5: Typical 3-Tier Network.....	226
Figure-6: Generic Service Template .....	235
Figure-7: Service template with network template A .....	236
Figure-8: Service template with network template B .....	236

---

# 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 XML specification ensuring that TOSCA semantics are preserved and can be transformed from XML to YAML or from YAML to XML. Additionally, in order to streamline the expression of TOSCA semantics, the YAML rendering is sought to be more concise and compact through the use of the YAML syntax.

## 1.2 Summary of key TOSCA concepts

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

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

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

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

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

## 63 1.4 Terminology

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

68  
69 The language provides an extension mechanism that can be used to extend the definitions with additional  
70 vendor-specific or domain-specific information.

## 71 1.5 Notational Conventions

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

### 75 1.5.1 Notes

- 76 • Sections that are titled “Example” throughout this document are considered non-normative.

## 77 1.6 Normative References

[RFC2119]	S. Bradner, <i>Key words for use in RFCs to Indicate Requirement Levels</i> , <a href="http://www.ietf.org/rfc/rfc2119.txt">http://www.ietf.org/rfc/rfc2119.txt</a> , IETF RFC 2119, March 1997.
[TOSCA-1.0]	Topology and Orchestration Topology and Orchestration Specification for Cloud Applications (TOSCA) Version 1.0, an OASIS Standard, 25 November 2013, <a href="http://docs.oasis-open.org/tosca/TOSCA/v1.0/os/TOSCA-v1.0-os.pdf">http://docs.oasis-open.org/tosca/TOSCA/v1.0/os/TOSCA-v1.0-os.pdf</a>
[YAML-1.2]	YAML, Version 1.2, 3rd Edition, Patched at 2009-10-01, Oren Ben-Kiki, Clark Evans, Ingy döt Net <a href="http://www.yaml.org/spec/1.2/spec.html">http://www.yaml.org/spec/1.2/spec.html</a>
[YAML-TS-1.1]	Timestamp Language-Independent Type for YAML Version 1.1, Working Draft 2005-01-18, <a href="http://yaml.org/type/timestamp.html">http://yaml.org/type/timestamp.html</a>

78 **1.7 Non-Normative References**

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

79 **1.8 Glossary**

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

<b>Term</b>	<b>Definition</b>
<b>Instance Model</b>	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.
<b>Node Template</b>	A <i>Node Template</i> specifies the occurrence of a software component node as part of a Topology Template. Each Node Template refers to a Node Type that defines the semantics of the node (e.g., properties, attributes, requirements, capabilities, interfaces). Node Types are defined separately for reuse purposes.
<b>Relationship Template</b>	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.
<b>Service Template</b>	A <i>Service Template</i> is typically used to specify the “topology” (or structure) and “orchestration” (or invocation of management behavior) of IT services so that they can be provisioned and managed in accordance with constraints and policies.  Specifically, TOSCA Service Templates optionally allow definitions of a TOSCA <a href="#">Topology Template</a> , TOSCA types (e.g., Node, Relationship, Capability, Artifact, etc.), groupings, policies and constraints along with any input or output declarations.
<b>Topology Model</b>	The term Topology Model is often used synonymously with the term <a href="#">Topology Template</a> with the use of “model” being prevalent when considering a Service Template’s topology definition as an <b>abstract representation</b> of an application or service to facilitate understanding of its functional components and by eliminating unnecessary details.
<b>Topology Template</b>	A Topology Template defines the structure of a service in the context of a Service Template. A Topology Template consists of a set of Node Template and Relationship Template definitions that together define the topology model of a service as a (not necessarily connected) directed graph.

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

---

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

---

---

## 82 2 TOSCA by example

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

### 86 2.1 A “hello world” template for TOSCA Simple Profile in YAML

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

#### 94 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
```

95 The template above contains a very simple topology template with only a single ‘Compute’ node template  
96 that declares some basic values for properties within two of the several capabilities that are built into the  
97 Compute node type definition. All TOSCA Orchestrators are expected to know how to instantiate a  
98 Compute node since it is normative and expected to represent a well-known function that is portable  
99 across TOSCA implementations. This expectation is true for all normative TOSCA Node and  
100 Relationship types that are defined in the Simple Profile specification. This means, with TOSCA’s

101 approach, that the application developer does not need to provide any deployment or implementation  
102 artifacts that contain code or logic to orchestrate these common software components. TOSCA  
103 orchestrators simply select or allocate the correct node (resource) type that fulfills the application  
104 topologies requirements using the properties declared in the node and its capabilities.

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

110  
111 The logical diagram of the “hello world” Compute node would look as follows:



112  
113  
114 As you can see, the **Compute** node also has attributes and other built-in capabilities, such as **Bindable**  
115 and **Endpoint**, each with additional properties that will be discussed in other examples later in this  
116 document. Although the Compute node has no direct properties apart from those in its capabilities, other  
117 TOSCA node type definitions may have properties that are part of the node type itself in addition to  
118 having Capabilities. TOSCA orchestration engines are expected to validate all property values provided  
119 in a node template against the property definitions in their respective node type definitions referenced in  
120 the service template. The `tosca_definitions_version` keyname in the TOSCA service template  
121 identifies the versioned set of normative TOSCA type definitions to use for validating those types defined  
122 in the TOSCA Simple Profile including the Compute node type. Specifically, the value  
123 `tosca_simple_yaml_1_0` indicates Simple Profile v1.0.0 definitions would be used for validation. Other  
124 type definitions may be imported from other service templates using the `import` keyword discussed later.

## 125 2.1.1 Requesting input parameters and providing output

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

### 131 Example 2 - Template with input and output parameter sections

```
tosca_definitions_version: tosca_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: toska.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 ] }

```

132 The **inputs** and **outputs** sections are contained in the **topology\_template** element of the TOSCA  
 133 template, meaning that they are scoped to node templates within the topology template. Input parameters  
 134 defined in the inputs section can be assigned to properties of node template within the containing  
 135 topology template; output parameters can be obtained from attributes of node templates within the  
 136 containing topology template.

137 Note that the **inputs** section of a TOSCA template allows for defining optional constraints on each input  
 138 parameter to restrict possible user input. Further note that TOSCA provides for a set of intrinsic functions  
 139 like **get\_input**, **get\_property** or **get\_attribute** to reference elements within the template or to  
 140 retrieve runtime values.

## 141 2.2 TOSCA template for a simple software installation

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

146 **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
```

147 The example above makes use of a node type **tosca.nodes.DBMS.MySQL** for the **mysql** node template to  
148 install MySQL on a server. This node type allows for setting a property **root\_password** to adapt the  
149 password of the MySQL root user at deployment. The set of properties and their schema has been  
150 defined in the node type definition. By means of the **get\_input** function, a value provided by the user at  
151 deployment time is used as value for the **root\_password** property. The same is true for the **port**  
152 property.

153 The **mysql** node template is related to the **db\_server** node template (of type **tosca.nodes.Compute**) via  
154 the **requirements** section to indicate where MySQL is to be installed. In the TOSCA metamodel, nodes  
155 get related to each other when one node has a requirement against some feature provided by another  
156 node. What kinds of requirements exist is defined by the respective node type. In case of MySQL, which  
157 is software that needs to be installed or hosted on a compute resource, the underlying node type named  
158 **DBMS** has a predefined requirement called **host**, which needs to be fulfilled by pointing to a node template  
159 of type **tosca.nodes.Compute**.

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



161

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

### 171 2.3 Overriding behavior of predefined node types

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

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

#### 181 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

```

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

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

## 190 2.4 TOSCA template for database content deployment

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

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

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

```


195 In the example above, the **my\_db** node template or type **tosca.nodes.Database.MySQL** represents an  
 196 actual MySQL database instance managed by a MySQL DBMS installation. The **requirements** section of  
 197 the **my\_db** node template expresses that the database it represents is to be hosted on a MySQL DBMS  
 198 node template named **mysql** which is also declared in this template.

199 In its **artifacts** section of the **my\_db** the node template, there is an artifact definition named **db\_content**  
 200 which represents a text file **my\_db\_content.txt** which in turn will be used to add content to the SQL

201 database as part of the **create** operation. The **requirements** section of the **my\_db** node template  
202 expresses that the database is hosted on a MySQL DBMS represented by the **mysql** node.

203 As you can see above, a script is associated with the create operation with the name **db\_create.sh**.  
204 The TOSCA Orchestrator sees that this is not a named artifact declared in the node's artifact section, but  
205 instead a filename for a normative TOSCA implementation artifact script type (i.e.,  
206 **tosca.artifacts.Implementation.Bash**). Since this is an implementation type for TOSCA, the  
207 orchestrator will execute the script automatically to create the node on **db\_server**, but first it will prepare  
208 the local environment with the declared inputs for the operation. In this case, the orchestrator would see  
209 that the **db\_data** input is using the **get\_artifact** function to retrieve the file (**my\_db\_content.txt**)  
210 which is associated with the **db\_content** artifact name prior to executing the **db\_create.sh** script.

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



212  
213 Note that while it would be possible to define one node type and corresponding node templates that  
214 represent both the DBMS middleware and actual database content as one entity, TOSCA normative node  
215 types distinguish between middleware (container) and application (containeer) node types. This allows on  
216 one hand to have better re-use of generic middleware node types without binding them to content running  
217 on top of them, and on the other hand this allows for better substitutability of, for example, middleware  
218 components like a DBMS during the deployment of TOSCA models.

## 219 2.5 TOSCA template for a two-tier application

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

226 **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: tosca.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: tosca.nodes.WebServer.Apache
  properties:
    # omitted here for brevity
  requirements:
    - host: web_server

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

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

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

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

```

227 The web application stack consists of the **wordpress** [WordPress], the **apache** [Apache] and the  
228 **web\_server** node templates. The **wordpress** node template represents a custom web application of type  
229 **tosca.nodes.WebApplication.WordPress** which is hosted on an Apache web server represented by the  
230 **apache** node template. This hosting relationship is expressed via the **host** entry in the **requirements**  
231 section of the **wordpress** node template. The **apache** node template, finally, is hosted on the  
232 **web\_server** compute node.

233 The database stack consists of the **wordpress\_db**, the **mysql** and the **db\_server** node templates. The  
234 **wordpress\_db** node represents a custom database of type **tosca.nodes.Database.MySQL** which is  
235 hosted on a MySQL DBMS represented by the **mysql** node template. This node, in turn, is hosted on the  
236 **db\_server** compute node.

237 The **wordpress** node requires a connection to the **wordpress\_db** node, since the WordPress application  
238 needs a database to store its data in. This relationship is established through the **database\_endpoint**  
239 entry in the **requirements** section of the **wordpress** node template's declared node type. For configuring  
240 the WordPress web application, information about the database to connect to is required as input to the  
241 **configure** operation. Therefore, the input parameters are defined and values for them are retrieved from  
242 the properties and attributes of the **wordpress\_db** node via the **get\_property** and **get\_attribute**  
243 functions. In the above example, these inputs are defined at the interface-level and would be available to  
244 all operations of the **Standard** interface (i.e., the **tosca.interfaces.node.lifecycle.Standard**  
245 interface) within the **wordpress** node template and not just the **configure** operation.

## 246 2.6 Using a custom script to establish a relationship in a template

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

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

### 258 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 ...

```

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

266 The **ConnectsTo** relationship (see 5.7.4) also provides a default **Configure** interface with operations that  
 267 optionally get executed when the orchestrator establishes the relationship. In the above example, the  
 268 author has provided the custom script **wp\_db\_configure.sh** to be executed for the operation called  
 269 **pre\_configure\_source**. The script file is assumed to be located relative to the referencing service  
 270 template such as a relative directory within the TOSCA Cloud Service Archive (CSAR) packaging format.  
 271 This approach allows for conveniently hooking in custom behavior without having to define a completely  
 272 new derived relationship type.

## 273 2.7 Using custom relationship types in a TOSCA template

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

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

## 280 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 ...
```

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

### 287 2.7.1 Definition of a custom relationship type

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

## 294 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
```

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

## 298 2.8 Defining generic dependencies between nodes in a template

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

## 303 Example 10 - Simple dependency relationship between two nodes

```
tosca_definitions_version: tosca_simple_yaml_1_0

description: Template with a generic dependency between two nodes.

topology_template:
  inputs:
    # omitted here for brevity

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

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

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

## 310 **2.9 Describing abstract requirements for nodes and capabilities in a** 311 **TOSCA template**

312 In TOSCA templates, nodes are either:

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

317

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

325

326 TOSCA supports two methods for template authors to express requirements for an abstract node within a  
327 TOSCA service template.

328

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

337

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

### 342 **2.9.1 Using a node\_filter to define hosting infrastructure requirements for a** 343 **software**

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

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

351 **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
```

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

357 Some of the constraints shown above narrow down the boundaries of allowed values for certain  
358 properties such as **mem\_size** or **num\_cpus** for the "**host**" capability by means of qualifier functions such  
359 as **greater\_or\_equal**. Other constraints, express specific values such as for the **architecture** or  
360 **distribution** properties of the "**os**" capability which will require the provider to find a precise match.

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

## 363 2.9.2 Using an abstract node template to define infrastructure requirements 364 for software

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

### 370 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
```

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

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

### 380 2.9.3 Using a `node_filter` to define requirements on a database for an 381 application

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

#### 386 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 }
```

387 In the example above, the application `my_app` requires a database node of type `MyDatabase` which has a  
388 `db_version` property value of `greater_or_equal` to the value 5.5.

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

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

398 **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 }
```

399 **2.10 Using node template substitution for model composition**

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

405 **2.10.1 Understanding node template instantiation through a TOSCA**  
406 **Orchestrator**

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

413  
414 In the latter case, a TOSCA Orchestrator will use additional substitution mapping information provided as  
415 part of the substituting topology templates to derive how the substituted part gets “wired” into the overall  
416 deployment, for example, how capabilities of a node template in the top-level topology template get  
417 bound to capabilities of node templates in the substituting topology template.


418

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

### 422 2.10.2 Definition of the top-level service template

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

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



432

433 **Figure 1: Using template substitution to implement a database tier**

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

441

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

446  
447 For example, if multiple components need to use the same database (or any other sub-system of the  
448 overall service), this can be expressed by means of normal relations between node templates, whereas  
449 such modeling would not be possible in **requirements** sections of disjoint node templates.

450 **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
```

451 **2.10.3 Definition of the database stack in a service template**


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

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

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

463 In short, the **substitution\_mappings** section provides the following information:

- 464 1. It defines what node templates, i.e. node templates of which type, can be substituted by the  
465 topology template.
- 466 2. It defines how capabilities of the substituted node (or the capabilities defined by the node type of  
467 the substituted node template, respectively) are bound to capabilities of node templates defined  
468 in the topology template.
- 469 3. It defines how requirements of the substituted node (or the requirements defined by the node type  
470 of the substituted node template, respectively) are bound to requirements of node templates  
471 defined in the topology template.



472  
473 **Figure 2: Substitution mappings**

474 The **substitution\_mappings** section in the sample below denotes that this topology template can be  
475 used for substituting node templates of type **tosca.nodes.Database**. It further denotes that the  
476 **database\_endpoint** capability of the substituted node gets fulfilled by the **database\_endpoint**  
477 capability of the **database** node contained in the topology template.

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

```
tosca_definitions_version: toska_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
```

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

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


## 487 **2.11 Using node template substitution for chaining subsystems**

488 A common use case when providing an end-to-end service is to define a chain of several subsystems that  
489 together implement the overall service. Those subsystems are typically defined as separate service  
490 templates to (1) keep the complexity of the end-to-end service template at a manageable level and to (2)

491 allow for the re-use of the respective subsystem templates in many different contexts. The type of  
 492 subsystems may be specific to the targeted workload, application domain, or custom use case. For  
 493 example, a company or a certain industry might define a subsystem type for company- or industry specific  
 494 data processing and then use that subsystem type for various end-user services. In addition, there might  
 495 be generic subsystem types like a database subsystem that are applicable to a wide range of use cases.

496 **2.11.1 Defining the overall subsystem chain**

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



506  
507

**Figure 3: Chaining of subsystems in a service template**

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

514

515 **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

```

516

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

521 Subsystems can be parameterized by providing properties. In the listing above, for example, the IP  
 522 address of the message queuing server is provided as property **mq\_service\_ip** to the transaction

523 processing subsystems and the desired port for receiving messages is specified by means of the  
524 **receiver\_port** property.

525 If attributes of the instantiated subsystems need to be obtained, this would be possible by using the  
526 **get\_attribute** intrinsic function on the respective subsystem node templates.

## 527 2.11.2 Defining a subsystem (node) type

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

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

534

### 535 Example 18 - Defining a TransactionSubsystem node type

```
tosca_definitions_version: tosca_simple_yaml_1_0

node_types:
  example.TransactionSubsystem:
    properties:
      mq_service_ip:
        type: string
      receiver_port:
        type: integer
    attributes:
      receiver_ip:
        type: string
      receiver_port:
        type: integer
    capabilities:
      message_receiver: tosca.capabilities.Endpoint
    requirements:
      - database_endpoint: tosca.capabilities.Endpoint.Database
```

536


537 Configuration parameters that would be allowed for customizing the instantiation of any subsystem are  
538 defined as properties of the node type. In the current example, those are the properties **mq\_service\_ip**  
539 and **receiver\_port** that had been used in the end-to-end service template in section 2.11.1.

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

## 543 2.11.3 Defining the details of a subsystem

544 The details of a subsystem, i.e. the software components and their hosting infrastructure, are defined as  
545 node templates and relationships in a service template. By means of substitution mappings that have  
546 been introduced in section 2.10.2, the service template is annotated to indicate to an orchestrator that it

547 can be used as realization of a node template of certain type, as well as how characteristics of the node  
 548 type are mapped to internal elements of the service template.  
 549



550  
 551 **Figure 4: Defining subsystem details in a service template**

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

555 The application named **app** provides a capability to receive messages, which is bound to the  
 556 **message\_receiver** capability of the substitutable node type. It further requires access to a database, so  
 557 the application's **database\_endpoint** requirement is mapped to the **database\_endpoint** requirement of  
 558 the **TransactionSubsystem** node type.

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

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

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

574

575 **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: tosca.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 ] }

```

## 576 2.12 Grouping node templates

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

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

### 585 Example 20 - Grouping Node Templates for possible policy application

```

tosca_definitions_version: toska_simple_yaml_1_0

description: Template for a scaling web server.

topology_template:
  inputs:
    # omitted here for brevity

  node_templates:
    apache:

```

```

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

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

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

```

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

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

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

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

#### 600 **Example 21 - Grouping nodes for anti-collocation policy application**

```

tosca_definitions_version: tosca_simple_yaml_1_0

description: Template hosting requirements and placement policy.

topology_template:
  inputs:
    # omitted here for brevity

  node_templates:
    wordpress_server:
      type: tosca.nodes.WebServer
      properties:

```

```

    # omitted here for brevity
requirements:
  - host:
    # Find a Compute node that fulfills these additional filter reqs.
    node_filter:
      capabilities:
        - host:
            properties:
              - mem_size: { greater_or_equal: 512 MB }
              - disk_size: { greater_or_equal: 2 GB }
        - os:
            properties:
              - architecture: x86_64
              - type: linux

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

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

policies:
  - my_anti_collocation_policy:
    type: my.policies.anticollocateion
    targets: [ my_co_location_group ]

```

```
# For this example, specific policy definitions are considered
# domain specific and are not included here
```

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

## 606 **2.13 Using YAML Macros to simplify templates**

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

613  
614 For example, an application that has a web server and database (i.e., a two-tier application) may be  
615 described using two **Compute** nodes (one to host the web server and another to host the database). The  
616 author may want both Compute nodes to be instantiated with similar properties such as operating system,  
617 distribution, version, etc.

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

### 621 **Example 22 - Using YAML anchors in TOSCA templates**

```
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

```

## 622 2.14 Passing information as inputs to Nodes and Relationships

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

### 629 2.14.1 Example: declaring input variables for all operations on a single 630 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 ] }

```

### 631 2.14.2 Example: declaring input variables for a single operation

```

node_templates:
  wordpress:
    type: tosca.nodes.WebApplication.WordPress
    requirements:
      ...
      - database_endpoint: mysql_database
    interfaces:
      Standard:
        create: wordpress_install.sh
        configure:
          implementation: wordpress_configure.sh
        inputs:

```

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

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

### 635 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
```

636  
637 In this example, the Standard create operation exposes / exports an environment variable named  
638 “generated\_url” attribute which will be assigned to the WordPress node’s url attribute.

### 639 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 ] }
```

640 In this example, the Standard lifecycle’s create operation exposes / exports an environment variable  
641 named “data\_dir” which will be passed as an input to the Standard lifecycle’s configure operation.

## 642 2.15 Topology Template Model versus Instance Model

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

## 653 2.16 Using attributes implicitly reflected from properties

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

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

662  
663 Use of this feature is shown in the example below where a source node named `my_client`, of type  
664 `ClientNode`, requires a connection to another node named `my_server` of type `ServerNode`. As you can  
665 see, the `ServerNode` type defines a property named `notification_port` which defines a dedicated port  
666 number which instances of `my_client` may use to post asynchronous notifications to it during runtime. In  
667 this case, the TOSCA Simple Profile assures that the `notification_port` property is implicitly reflected  
668 as an attribute in the `my_server` node (also with the name `notification_port`) when its node template  
669 is instantiated.

670

### 671 Example 23 - Properties reflected as attributes

```
tosca_definitions_version: tosca_simple_yaml_1_0

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

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

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

```

topology_template:
  node_templates:

    my_server:
      type: ServerNode
      properties:
        notification_port: 8000

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

  relationship_templates:
    my_connection:
      type: ConnectsTo
      interfaces:
        Configure:
          inputs:
            targ_notify_port: { get_attribute: [ TARGET, notification_port ] }
            # other operation definitions omitted here for brevity

```

672

673 Specifically, the above example shows that the **ClientNode** type needs the **notification\_port** value  
 674 anytime a node of **ServerType** is connected to it using the **ConnectsTo** relationship in order to make it  
 675 available to its **Configure** operations (scripts). It does this by using the **get\_attribute** function to  
 676 retrieve the **notification\_port** attribute from the **TARGET** node of the **ConnectsTo** relationship (which is  
 677 a node of type **ServerNode**) and assigning it to an environment variable named **targ\_notify\_port**.

678

679 It should be noted that the actual port value of the **notification\_port** attribute may or may not be the  
 680 value **8000** as requested on the property; therefore, any node that is dependent on knowing its actual  
 681 “runtime” value would use the **get\_attribute** function instead of the **get\_property** function.

682

## 3 TOSCA Simple Profile definitions in YAML

683  
684  
685

Except for the examples, 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).

686

### 3.1 TOSCA Namespace URI and alias

687  
688

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_1	<a href="http://docs.oasis-open.org/tosca/ns/simple/yaml/1.1">http://docs.oasis-open.org/tosca/ns/simple/yaml/1.1</a>	The TOSCA Simple Profile v1.1 (YAML) target namespace and namespace alias.

689

#### 3.1.1 TOSCA Namespace prefix

690  
691

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

692

#### 3.1.2 TOSCA Namespacing in TOSCA Service Templates

693  
694  
695

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:

696  
697  
698  
699  
700  
701  
702  
703  
704

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.

705

#### 3.1.3 Rules to avoid namespace collisions

706  
707  
708  
709  
710  
711

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.

712

##### 3.1.3.1 Additional Requirements

713  
714  
715

- 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:

- 716 ○ Imported type definitions with the same Namespace URI, local name and version SHALL
- 717 be equivalent.
- 718 ○ If different values of the “`tosca_definitions_version`” are encountered, their
- 719 corresponding type definitions MUST be uniquely identifiable using their corresponding
- 720 Namespace URI using a different Namespace prefix.
- 721 ● Duplicate local names (i.e., within the same Service Template SHALL be considered an error.
- 722 These include, but are not limited to duplicate names found for the following definitions:
- 723 ○ Repositories (`repositories`)
- 724 ○ Data Types (`data_types`)
- 725 ○ Node Types (`node_types`)
- 726 ○ Relationship Types (`relationship_types`)
- 727 ○ Capability Types (`capability_types`)
- 728 ○ Artifact Types (`artifact_types`)
- 729 ○ Interface Types (`interface_types`)
- 730 ● Duplicate Template names within a Service Template’s Topology Template SHALL be considered
- 731 an error. These include, but are not limited to duplicate names found for the following template
- 732 types:
- 733 ○ Node Templates (`node_templates`)
- 734 ○ Relationship Templates (`relationship_templates`)
- 735 ○ Inputs (`inputs`)
- 736 ○ Outputs (`outputs`)
- 737 ○ Groups (`groups`)
- 738 ● Duplicate names for the following keynames within Types or Templates SHALL be considered an
- 739 error. These include, but are not limited to duplicate names found for the following keynames:
- 740 ○ Properties (`properties`)
- 741 ○ Attributes (`attributes`)
- 742 ○ Artifacts (`artifacts`)
- 743 ○ Requirements (`requirements`)
- 744 ○ Capabilities (`capabilities`)
- 745 ○ Interfaces (`interfaces`)

## 746 3.2 Parameter and property types

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

### 749 3.2.1 Referenced YAML Types

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

752 The following table declares the valid YAML type URIs and aliases that SHALL be used when possible  
753 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	<a href="#">tag:yaml.org,2002:timestamp</a> [ <a href="#">YAML-TS-1.1</a> ]

null	tag:yaml.org,2002:null
------	------------------------

### 754 3.2.1.1 Notes

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

### 758 3.2.2 TOSCA version

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

<b>Shorthand Name</b>	version
<b>Type Qualified Name</b>	tosca:version

### 764 3.2.2.1 Grammar

765 TOSCA version strings have the following grammar:

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

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

- 767 • **major\_version**: is a required integer value greater than or equal to 0 (zero)
- 768 • **minor\_version**: is a required integer value greater than or equal to 0 (zero).
- 769 • **fix\_version**: is an optional integer value greater than or equal to 0 (zero).
- 770 • **qualifier**: is an optional string that indicates a named, pre-release version of the associated
- 771 code that has been derived from the version of the code identified by the combination
- 772 **major\_version**, **minor\_version** and **fix\_version** numbers.
- 773 • **build\_version**: is an optional integer value greater than or equal to 0 (zero) that can be used to
- 774 further qualify different build versions of the code that has the same **qualifer\_string**.

### 775 3.2.2.2 Version Comparison

- 776 • When comparing TOSCA versions, all component versions (i.e., *major*, *minor* and *fix*) are
- 777 compared in sequence from left to right.
- 778 • TOSCA versions that include the optional qualifier are considered older than those without a
- 779 qualifier.
- 780 • TOSCA versions with the same major, minor, and fix versions and have the same qualifier string,
- 781 but with different build versions can be compared based upon the build version.
- 782 • Qualifier strings are considered domain-specific. Therefore, this specification makes no
- 783 recommendation on how to compare TOSCA versions with the same major, minor and fix
- 784 versions, but with different qualifiers strings and simply considers them different named branches
- 785 derived from the same code.

### 786 3.2.2.3 Examples

787 Examples of valid TOSCA version strings:

```
# basic version strings
6.1
2.0.1

# version string with optional qualifier
3.1.0.beta

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

### 788 3.2.2.4 Notes

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

### 791 3.2.2.5 Additional Requirements

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

## 794 3.2.3 TOSCA range type

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

<b>Shorthand Name</b>	range
<b>Type Qualified Name</b>	tosca:range

### 797 3.2.3.1 Grammar

798 TOSCA range values have the following grammar:

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

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

- 800 • **lower\_bound**: is a required integer value that denotes the lower boundary of the range.
- 801 • **upper\_bound**: is a required integer value that denotes the upper boundary of the range. This  
802 value MUST be greater than **lower\_bound**.

### 803 3.2.3.2 Keywords

804 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.

### 805 3.2.3.3 Examples

806 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 ]
```

807

## 808 3.2.4 TOSCA list type

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

812 Note that entries in a list for one property or parameter must be of the same type. The type (for simple  
813 entries) or schema (for complex entries) is defined by the **entry\_schema** attribute of the respective  
814 [property definition](#), [attribute definitions](#), or input or output [parameter definitions](#).

<b>Shorthand Name</b>	list
<b>Type Qualified Name</b>	tosca:list

### 815 3.2.4.1 Grammar

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

#### 817 3.2.4.1.1 Square bracket notation

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

#### 818 3.2.4.1.2 Bulleted (sequenced) list notation

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

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

- 820 • **<list\_entry\_\*>**: represents one entry of the list.

### 821 3.2.4.2 Declaration Examples

#### 822 3.2.4.2.1 List declaration using a simple type

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

```
<some_entity>:
  ...
  properties:
    listen_ports:
```

```
type: list
entry_schema:
  description: listen port entry (simple integer type)
  type: integer
  constraints:
    - max_length: 128
```

### 825 3.2.4.2.2 List declaration using a complex type

826 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
```

### 827 3.2.4.3 Definition Examples

828 These examples show two notation options for defining lists:

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

#### 832 3.2.4.3.1 Square bracket notation

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

#### 833 3.2.4.3.2 Bulleted list notation

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

### 834 3.2.5 TOSCA map type

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

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

<b>Shorthand Name</b>	map
<b>Type Qualified Name</b>	tosca:map

### 841 3.2.5.1 Grammar

842 TOSCA maps are normal YAML dictionaries with following grammar:

#### 843 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>
```

#### 844 3.2.5.1.2 Multi-line grammar

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

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

- 846 • **entry\_key\_\***: is the required key for an entry in the map
- 847 • **entry\_value\_\***: is the value of the respective entry in the map

### 848 3.2.5.2 Declaration Examples

#### 849 3.2.5.2.1 Map declaration using a simple type

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

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

#### 852 3.2.5.2.2 Map declaration using a complex type

853 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
```

### 854 3.2.5.3 Definition Examples

855 These examples show two notation options for defining maps:

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

#### 859 3.2.5.3.1 Single-line notation

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

#### 860 3.2.5.3.2 Multi-line notation

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

### 861 3.2.6 TOSCA scalar-unit type

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

#### 864 3.2.6.1 Grammar

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

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

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

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

#### 869 3.2.6.2 Additional requirements

- 870 • **Whitespace**: any number of spaces (including zero or none) **SHALL** be allowed between the  
871 **scalar** value and the **unit** value.
- 872 • It **SHALL** be considered an error if either the scalar or unit portion is missing on a property or  
873 attribute declaration derived from any scalar-unit type.

- 874
- 875
- 876
- 877
- 878
- When performing constraint clause evaluation on values of the scalar-unit type, both the scalar value portion and unit value portion **SHALL** be compared together (i.e., both are treated as a single value). For example, if we have a property called **storage\_size**. which is of type scalar-unit, a valid range constraint would appear as follows:
    - `storage_size: in_range [ 4 GB, 20 GB ]`

879 where **storage\_size**'s range would be evaluated using both the numeric and unit values  
 880 (combined together), in this case '4 GB' and '20 GB'.

881 **3.2.6.3 Concrete Types**

<b>Shorthand Names</b>	scalar-unit.size, scalar-unit.size
<b>Type Qualified Names</b>	tosca:scalar-unit.size, toasca:scalar-unit.time

- 882
- 883 The scalar-unit type grammar is abstract and has two recognized concrete types in TOSCA:
- 884
- 885
- 886
- 887
- **scalar-unit.size** – used to define properties that have scalar values measured in size units.
  - **scalar-unit.time** – used to define properties that have scalar values measured in size units.
  - **scalar-unit.frequency** – used to define properties that have scalar values measured in units per second.

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

889 **3.2.6.4 scalar-unit.size**

890 **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)

891 **3.2.6.4.2 Examples**

```
# Storage size in Gigabytes
properties:
```

```
storage_size: 10 GB
```

### 892 3.2.6.4.3 Notes

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

### 903 3.2.6.5 scalar-unit.time

#### 904 3.2.6.5.1 Recognized Units

Unit	Usage	Description
d	time	days
h	time	hours
m	time	minutes
s	time	seconds
ms	time	milliseconds
us	time	microseconds
ns	time	nanoseconds

#### 905 3.2.6.5.2 Examples

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

#### 906 3.2.6.5.3 Notes

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

914 **3.2.6.6 scalar-unit.frequency**

915 **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.

916 **3.2.6.6.2 Examples**

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

917 **3.2.6.6.3 Notes**

- 918
- The value for Hertz (Hz) is the International Standard Unit (ISU) as described by the Bureau International des Poids et Mesures (BIPM) in the “*SI Brochure: The International System of Units (SI) [8th edition, 2006; updated in 2014]*”, <http://www.bipm.org/en/publications/si-brochure/>
- 919
- 920

921 **3.3 Normative values**

922 **3.3.1 Node States**

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

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

Node State		
Value	Transitional	Description
initial	no	Node is not yet created. Node only exists as a template definition.
creating	yes	Node is transitioning from <b>initial</b> state to <b>created</b> state.
created	no	Node software has been installed.
configuring	yes	Node is transitioning from <b>created</b> state to <b>configured</b> state.
configured	no	Node has been configured prior to being started.
starting	yes	Node is transitioning from <b>configured</b> state to <b>started</b> state.

Node State		
Value	Transitional	Description
started	no	Node is started.
stopping	yes	Node is transitioning from its current state to a <b>configured</b> 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.

### 929 3.3.2 Relationship States

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

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

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

#### 934 3.3.2.1 Notes

- 935 • Additional states may be defined in future versions of the TOSCA Simple Profile in YAML  
936 specification.

### 937 3.3.3 Directives

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

### 939 3.3.4 Network Name aliases

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

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

### 943 3.3.4.1 Usage

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

## 947 3.4 TOSCA Metamodel

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

### 950 3.4.1 Required Keynames

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

## 956 3.5 Reusable modeling definitions

### 957 3.5.1 Description definition

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

#### 960 3.5.1.1 Keyname

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

```
description
```

#### 962 3.5.1.2 Grammar

963 Description definitions have the following grammar:

```
description: <string>
```

#### 964 3.5.1.3 Examples

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

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

967 The YAML “folded” style may also be used for multi-line descriptions which “folds” line breaks as space  
968 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.

#### 969 3.5.1.4 Notes

- 970
- Use of “folded” style is discouraged for the YAML string type apart from when used with the **description** keyname.
- 971

#### 972 3.5.2 Constraint clause

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

#### 976 3.5.2.1 Operator keynames

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

Operator	Type	Value Type	Description
equal	scalar	any	Constrains a property or parameter to a value equal to ('=') the value declared.
greater_than	scalar	comparable	Constrains a property or parameter to a value greater than ('>') the value declared.
greater_or_equal	scalar	comparable	Constrains a property or parameter to a value greater than or equal to ('>=') the value declared.
less_than	scalar	comparable	Constrains a property or parameter to a value less than ('<') the value declared.
less_or_equal	scalar	comparable	Constrains a property or parameter to a value less than or equal to ('<=') the value declared.
in_range	dual scalar	comparable, <a href="#">range</a>	Constrains a property or parameter to a value in range of (inclusive) the two values declared.  <b>Note:</b> subclasses or templates of types that declare a property with the <b>in_range</b> 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	<a href="#">string</a> , <a href="#">list</a> , <a href="#">map</a>	Constrains the property or parameter to a value of a given length.
min_length	scalar	<a href="#">string</a> , <a href="#">list</a> , <a href="#">map</a>	Constrains the property or parameter to a value to a minimum length.
max_length	scalar	<a href="#">string</a> , <a href="#">list</a> , <a href="#">map</a>	Constrains the property or parameter to a value to a maximum length.
pattern	regex	<a href="#">string</a>	Constrains the property or parameter to a value that is allowed by the provided regular expression.  <b>Note:</b> Future drafts of this specification will detail the use of regular expressions and reference an appropriate standardized grammar.

### 978 3.5.2.1.1 Comparable value types

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

### 982 3.5.2.2 Additional Requirements

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

### 991 3.5.2.3 Grammar

992 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>
```

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

- 994 • **operator**: represents a required operator from the specified list shown above (see section  
995 3.5.2.1 “Operator keynames”).
- 996 • **scalar\_value**, **scalar\_value\_\***: represents a required scalar (or atomic quantity) that can  
997 hold only one value at a time. This will be a value of a primitive type, such as an integer or string  
998 that is allowed by this specification.
- 999 • **value\_\***: represents a required value of the operator that is not limited to scalars.
- 1000 • **regular\_expression\_value**: represents a regular expression (string) value.

### 1001 3.5.2.4 Examples

1002 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

```

### 1003 **3.5.3 Property Filter definition**

1004 A property filter definition defines criteria, using constraint clauses, for selection of a TOSCA entity based  
 1005 upon its property values.

#### 1006 **3.5.3.1 Grammar**

1007 Property filter definitions have one of the following grammars:

##### 1008 **3.5.3.1.1 Short notation:**

1009 The following single-line grammar may be used when only a single constraint is needed on a property:

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

##### 1010 **3.5.3.1.2 Extended notation:**

1011 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>
```

1012 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1013 • **property\_name**: represents the name of property that would be used to select a property
- 1014 definition with the same name (**property\_name**) on a TOSCA entity (e.g., a Node Type, Node
- 1015 Template, Capability Type, etc.).
- 1016 • **property\_constraint\_clause\_\***: represents constraint clause(s) that would be used to filter
- 1017 entities based upon the named property's value(s).

### 1018 3.5.3.2 Additional Requirements

- 1019 • Property constraint clauses must be type compatible with the property definitions (of the same
- 1020 name) as defined on the target TOSCA entity that the clause would be applied against.

### 1021 3.5.4 Node Filter definition

1022 A node filter definition defines criteria for selection of a TOSCA Node Template based upon the  
1023 template's property values, capabilities and capability properties.

#### 1024 3.5.4.1 Keynames

1025 The following is the list of recognized keynames for a TOSCA node filter definition:

Keyname	Required	Type	Description
properties	no	list of <a href="#">property filter definition</a>	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 <a href="#">capability type names</a>	An optional sequenced list of capability names or types that would be used to select (filter) matching TOSCA entities based upon their existence.

#### 1026 3.5.4.2 Additional filtering on named Capability properties

1027 Capabilities used as filters often have their own sets of properties which also can be used to construct a  
1028 filter.

Keyname	Required	Type	Description
<capability name_or_type> name>: properties	no	list of <a href="#">property filter definitions</a>	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.

#### 1029 3.5.4.3 Grammar

1030 Node filter definitions have following grammar:

```
<filter_name>:  
  properties:
```

```

- <property_filter_def 1>
- ...
- <property_filter_def n>
capabilities:
- <capability_name_or_type_1>:
  properties:
    - <cap 1 property_filter_def 1>
    - ...
    - <cap m property_filter_def n>
- ...
- <capability_name_or_type_n>:
  properties:
    - <cap 1 property_filter_def 1>
    - ...
    - <cap m property_filter_def n>

```

1031 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1032 • **property\_filter\_def\_\***: represents a property filter definition that would be used to select
- 1033 (filter) matching TOSCA entities (e.g., Node Template, Node Type, Capability Types, etc.) based
- 1034 upon their property definitions' values.
- 1035 • **capability\_name\_or\_type\_\***: represents the type or name of a capability that would be used
- 1036 to select (filter) matching TOSCA entities based upon their existence.
- 1037 • **cap\_\*\_property\_def\_\***: represents a property filter definition that would be used to select
- 1038 (filter) matching TOSCA entities (e.g., Node Template, Node Type, Capability Types, etc.) based
- 1039 upon their capabilities' property definitions' values.

#### 1040 3.5.4.4 Additional requirements

- 1041 • TOSCA orchestrators **SHALL** search for matching capabilities listed on a target filter by assuming
- 1042 the capability name is first a symbolic name and secondly it is a type name (in order to avoid
- 1043 namespace collisions).

#### 1044 3.5.4.5 Example

1045 The following example is a filter that would be used to select a TOSCA [Compute](#) node based upon the

1046 values of its defined capabilities. Specifically, this filter would select Compute nodes that supported a

1047 specific range of CPUs (i.e., `num_cpus` value between 1 and 4) and memory size (i.e., `mem_size` of 2 or

1048 greater) from its declared "host" capability.

1049

```

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 }

```

## 1050 3.5.5 Repository definition

1051 A repository definition defines a named external repository which contains deployment and  
1052 implementation artifacts that are referenced within the TOSCA Service Template.

### 1053 3.5.5.1 Keynames

1054 The following is the list of recognized keynames for a TOSCA repository definition:

Keyname	Required	Type	Constraints	Description
description	no	<a href="#">description</a>	None	The optional description for the repository.
url	yes	<a href="#">string</a>	None	The required URL or network address used to access the repository.
credential	no	<a href="#">Credential</a>	None	The optional Credential used to authorize access to the repository.

### 1055 3.5.5.2 Grammar

1056 Repository definitions have one the following grammars:

#### 1057 3.5.5.2.1 Single-line grammar (no credential):

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

#### 1058 3.5.5.2.2 Multi-line grammar

```

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

```

1059 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1060 • **repository\_name**: represents the required symbolic name of the repository as a [string](#).
- 1061 • **repository\_description**: contains an optional description of the repository.
- 1062 • **repository\_address**: represents the required URL of the repository as a string.
- 1063 • **authorization\_credential**: represents the optional credentials (e.g., user ID and password)  
1064 used to authorize access to the repository.

### 1065 3.5.5.3 Example

1066 The following represents a repository definition:

```

repositories:
  my_code_repo:

```

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

### 1067 3.5.6 Artifact definition

1068 An artifact definition defines a named, typed file that can be associated with Node Type or Node  
1069 Template and used by orchestration engine to facilitate deployment and implementation of interface  
1070 operations.

#### 1071 3.5.6.1 Keynames

1072 The following is the list of recognized keynames for a TOSCA artifact definition when using the extended  
1073 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 <b>file</b> 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.

#### 1074 3.5.6.2 Grammar

1075 Artifact definitions have one of the following grammars:

##### 1076 3.5.6.2.1 Short notation

1077 The following single-line grammar may be used when the artifact's type and mime type can be inferred  
1078 from the file URI:

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

##### 1079 3.5.6.2.2 Extended notation:

1080 The following multi-line grammar may be used when the artifact's definition's type and mime type need to  
1081 be explicitly declared:

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

1082 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1083 • **artifact\_name**: represents the required symbolic name of the artifact as a **string**.

- 1084 • **artifact\_description**: represents the optional [description](#) for the artifact.
- 1085 • **artifact\_type\_name**: represents the required [artifact type](#) the artifact definition is based upon.
- 1086 • **artifact\_file\_URI**: **represents the required URI string (relative or absolute) which can be**
- 1087 **used to locate the artifact's file.**
- 1088 • **artifact\_repository\_name**: represents the optional name of the [repository definition](#) to use to
- 1089 retrieve the associated artifact (file) from.
- 1090 • **file\_deployment\_path**: represents the optional path the **artifact\_file\_URI** would be
- 1091 copied into within the target node's container.

### 1092 3.5.6.3 Example

1093 The following represents an artifact definition:

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

### 1094 3.5.7 Import definition

1095 An import definition is used within a TOSCA Service Template to locate and uniquely name another  
 1096 TOSCA Service Template file which has type and template definitions to be imported (included) and  
 1097 referenced within another Service Template.

#### 1098 3.5.7.1 Keynames

1099 The following is the list of recognized keynames for a TOSCA import definition:

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

#### 1100 3.5.7.2 Grammar

1101 Import definitions have one the following grammars:

##### 1102 3.5.7.2.1 Single-line grammar:

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

##### 1103 3.5.7.2.2 Multi-line grammar

```
imports:
  - file: <file_URI>
```

```

repository: <repository_name>
namespace_uri: <definition_namespace_uri>
namespace_prefix: <definition_namespace_prefix>

```

1104 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1105 • **file\_uri**: contains the required name (i.e., URI) of the file to be imported as a [string](#).
- 1106 • **repository\_name**: represents the optional symbolic name of the repository definition where the
- 1107 imported file can be found as a [string](#).
- 1108 • **namespace\_uri**: represents the optional namespace URI to that will be applied to type
- 1109 definitions found within the imported file as a [string](#).
- 1110 • **namespace\_prefix**: represents the optional namespace prefix (alias) that will be used to
- 1111 indicate the **namespace\_uri** when forming a qualified name (i.e., QName) when referencing type
- 1112 definitions from the imported file as a [string](#).

### 1113 3.5.7.3 Example

1114 The following represents how import definitions would be used for the imports keyname within a TOSCA  
1115 Service Template:

```

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

```

## 1116 3.5.8 Property definition

1117 A property definition defines a named, typed value and related data that can be associated with an entity  
1118 defined in this specification (e.g., Node Types, Relationship Types, Capability Types, etc.). Properties  
1119 are used by template authors to provide input values to TOSCA entities which indicate their “desired  
1120 state” when they are instantiated. The value of a property can be retrieved using the **get\_property**  
1121 function within TOSCA Service Templates.

### 1122 3.5.8.1.1 Attribute and Property reflection

1123 The actual state of the entity, at any point in its lifecycle once instantiated, is reflected by [Attribute](#)  
1124 [definitions](#). TOSCA orchestrators automatically create an attribute for every declared property (with the  
1125 same symbolic name) to allow introspection of both the desired state (property) and actual state  
1126 (attribute).

### 1127 3.5.8.2 Keynames

1128 The following is the list of recognized keynames for a TOSCA property definition:

Keyname	Required	Type	Constraints	Description
type	yes	<a href="#">string</a>	None	The required data type for the property.
description	no	<a href="#">description</a>	None	The optional description for the property.
required	no	<a href="#">boolean</a>	default: true	An optional key that declares a property as required ( <b>true</b> ) or not ( <b>false</b> ).

Keyname	Required	Type	Constraints	Description
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	<a href="#">string</a>	default: supported	The optional status of the property relative to the specification or implementation. See table below for valid values.
constraints	no	list of <a href="#">constraint clauses</a>	None	The optional list of sequenced constraint clauses for the property.
entry_schema	no	<a href="#">string</a>	None	The optional key that is used to declare the name of the <a href="#">Datatype definition</a> for entries of set types such as the TOSCA <a href="#">list</a> or <a href="#">map</a> .

1129 **3.5.8.3 Status values**

1130 The following property status values are supported:

Value	Description
<b>supported</b>	Indicates the property is supported. This is the <b>default</b> value for all property definitions.
<b>unsupported</b>	Indicates the property is not supported.
<b>experimental</b>	Indicates the property is experimental and has no official standing.
<b>deprecated</b>	Indicates the property has been deprecated by a new specification version.

1131 **3.5.8.4 Grammar**

1132 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>

```

1133 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1134 • **property\_name**: represents the required symbolic name of the property as a [string](#).
- 1135 • **property\_description**: represents the optional [description](#) of the property.
- 1136 • **property\_type**: represents the required data type of the property.

- 1137 • **property\_required**: represents an optional [boolean](#) value (true or false) indicating whether or  
1138 not the property is required. If this keyname is not present on a property definition, then the  
1139 property SHALL be considered **required** (i.e., true) by **default**.
- 1140 • **default\_value**: contains a type-compatible value that may be used as a default if not provided  
1141 by another means.
- 1142 • **status\_value**: a [string](#) that contains a keyword that indicates the status of the property relative  
1143 to the specification or implementation.
- 1144 • **property\_constraints**: represents the optional [sequenced](#) list of one or more [constraint](#)  
1145 [clauses](#) on the property definition.
- 1146 • **entry\_description**: represents the optional [description](#) of the entry schema.
- 1147 • **entry\_type**: represents the required type name for entries in a [list](#) or [map](#) property type.
- 1148 • **entry\_constraints**: represents the optional [sequenced](#) list of one or more [constraint clauses](#)  
1149 on entries in a [list](#) or [map](#) property type.

### 1150 3.5.8.5 Additional Requirements

- 1151 • Implementations of the TOSCA Simple Profile **SHALL** automatically reflect (i.e., make available)  
1152 any property defined on an entity as an attribute of the entity with the same name as the property.
- 1153 • A property **SHALL** be considered [required by default](#) (i.e., as if the **required** keyname on the  
1154 definition is set to **true**) unless the definition's **required** keyname is explicitly set to **false**.
- 1155 • The value provided on a property definition's **default** keyname **SHALL** be type compatible with  
1156 the type declared on the definition's **type** keyname.
- 1157 • Constraints of a property definition **SHALL** be type-compatible with the type defined for that  
1158 definition.

### 1159 3.5.8.6 Notes

- 1160 • This element directly maps to the **PropertiesDefinition** element defined as part of the  
1161 schema for most type and entities defined in the [TOSCA v1.0 specification](#).
- 1162 • In the [TOSCA v1.0 specification](#) constraints are expressed in the XML Schema definitions of  
1163 Node Type properties referenced in the **PropertiesDefinition** element of **NodeType**  
1164 definitions.

### 1165 3.5.8.7 Example

1166 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 ]
```

### 1167 3.5.9 Property assignment

1168 This section defines the grammar for assigning values to named properties within TOSCA Node and  
1169 Relationship templates that are defined in their corresponding named types.

1170 **3.5.9.1 Keynames**

1171 The TOSCA property assignment has no keynames.

1172 **3.5.9.2 Grammar**

1173 Property assignments have the following grammar:

1174 **3.5.9.2.1 Short notation:**

1175 The following single-line grammar may be used when a simple value assignment is needed:

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

1176 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1177 • **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.).
- 1178
- 1179
- 1180
- 1181 • **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
- 1182 evaluation of an expression or a function.
- 1183

1184 **3.5.10 Attribute definition**

1185 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.

1190 **3.5.10.1 Attribute and Property reflection**

1191 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.

1194 **3.5.10.2 Keynames**

1195 The following is the list of recognized keynames for a TOSCA attribute definition:

Keyname	Required	Type	Constraints	Description
type	yes	<a href="#">string</a>	None	The required data type for the attribute.
description	no	<a href="#">description</a>	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 <b>type</b> keyname.
status	no	<a href="#">string</a>	default: <a href="#">supported</a>	The optional status of the attribute relative to the specification or implementation. See supported <a href="#">status values</a> defined under the <a href="#">Property definition</a> section.
entry_schema	no	<a href="#">string</a>	None	The optional key that is used to declare the name of the <a href="#">Datatype definition</a> for entries of set types such as the TOSCA <a href="#">list</a> or <a href="#">map</a> .

### 1196 3.5.10.3 Grammar

1197 Attribute definitions have the following grammar:

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

1198 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1199 • **attribute\_name**: represents the required symbolic name of the attribute as a [string](#).
- 1200 • **attribute\_type**: represents the required data type of the attribute.
- 1201 • **attribute\_description**: represents the optional [description](#) of the attribute.
- 1202 • **default\_value**: contains a type-compatible value that may be used as a default if not provided
- 1203 by another means.
- 1204 • **status\_value**: contains a value indicating the attribute's status relative to the specification
- 1205 version (e.g., supported, deprecated, etc.). Supported [status values](#) for this keyname are defined
- 1206 under [Property definition](#).

### 1207 3.5.10.4 Additional Requirements

- 1208 • In addition to any explicitly defined attributes on a TOSCA entity (e.g., Node Type,
- 1209 RelationshipType, etc.), implementations of the TOSCA Simple Profile **MUST** automatically
- 1210 reflect (i.e., make available) any property defined on an entity as an attribute of the entity with the
- 1211 same name as the property.
- 1212 • Values for the default keyname **MUST** be derived or calculated from other attribute or operation
- 1213 output values (that reflect the actual state of the instance of the corresponding resource) and not
- 1214 hard-coded or derived from a property settings or inputs (i.e., desired state).

### 1215 3.5.10.5 Notes

- 1216 • Attribute definitions are very similar to [Property definitions](#); however, properties of entities reflect
- 1217 an input that carries the template author's requested or desired value (i.e., desired state) which
- 1218 the orchestrator (attempts to) use when instantiating the entity whereas attributes reflect the
- 1219 actual value (i.e., actual state) that provides the actual instantiated value.
  - 1220 ○ For example, a property can be used to request the IP address of a node using a
  - 1221 property (setting); however, the actual IP address after the node is instantiated may by
  - 1222 different and made available by an attribute.

### 1223 3.5.10.6 Example

1224 The following represents a required attribute definition:

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

## 1225 3.5.11 Attribute assignment

1226 This section defines the grammar for assigning values to named attributes within TOSCA Node and  
1227 Relationship templates which are defined in their corresponding named types.

### 1228 3.5.11.1 Keynames

1229 The TOSCA attribute assignment has no keynames.

### 1230 3.5.11.2 Grammar

1231 Attribute assignments have the following grammar:

#### 1232 3.5.11.2.1 Short notation:

1233 The following single-line grammar may be used when a simple value assignment is needed:

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

#### 1234 3.5.11.2.2 Extended notation:

1235 The following multi-line grammar may be used when a value assignment requires keys in addition to a  
1236 simple value assignment:

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

1237 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1238 • **attribute\_name**: represents the name of an attribute that would be used to select an attribute  
1239 definition with the same name within on a TOSCA entity (e.g., Node Template, Relationship  
1240 Template, etc.) which is declared (or reflected from a Property definition) in its declared type  
1241 (e.g., a Node Type, Node Template, Capability Type, etc.).
- 1242 • **attribute\_value**, **attribute\_value\_expression**: represent the type-compatible value to  
1243 assign to the named attribute. Attribute values may be provided as the result from the  
1244 evaluation of an expression or a function.
- 1245 • **attribute\_description**: represents the optional [description](#) of the attribute.

### 1246 3.5.11.3 Additional requirements

- 1247 • Attribute values **MAY** be provided by the underlying implementation at runtime when requested  
1248 by the `get_attribute` function or it **MAY** be provided through the evaluation of expressions and/or  
1249 functions that derive the values from other TOSCA attributes (also at runtime).

## 1250 3.5.12 Parameter definition

1251 A parameter definition is essentially a TOSCA property definition; however, it also allows a value to be  
1252 assigned to it (as for a TOSCA property assignment). In addition, in the case of output parameters, it can  
1253 optionally inherit the data type of the value assigned to it rather than have an explicit data type defined for  
1254 it.

### 1255 3.5.12.1 Keynames

1256 The TOSCA parameter definition has all the keynames of a TOSCA Property definition, but in addition  
1257 includes the following additional or changed keynames:

Keyname	Required	Type	Constraints	Description
type	no	<a href="#">string</a>	None	The required data type for the parameter.  <b>Note:</b> 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.

1258 **3.5.12.2 Grammar**

1259 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>

```

1260 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1261 • **parameter\_name**: represents the required symbolic name of the parameter as a [string](#).
- 1262 • **parameter\_description**: represents the optional [description](#) of the parameter.
- 1263 • **parameter\_type**: represents the optional data type of the parameter. Note, this keyname is
- 1264 required for a TOSCA Property definition, but is not for a TOSCA Parameter definition.
- 1265 • **parameter\_value**, **parameter\_value\_expresssion**: represent the type-compatible value to
- 1266 assign to the named parameter. Parameter values may be provided as the result from the
- 1267 evaluation of an expression or a function.
- 1268 • **parameter\_required**: represents an optional [boolean](#) value (true or false) indicating whether or
- 1269 not the parameter is required. If this keyname is not present on a parameter definition, then the
- 1270 property SHALL be considered **required** (i.e., true) by **default**.
- 1271 • **default\_value**: contains a type-compatible value that may be used as a default if not provided
- 1272 by another means.
- 1273 • **status\_value**: a [string](#) that contains a keyword that indicates the status of the parameter
- 1274 relative to the specification or implementation.
- 1275 • **parameter\_constraints**: represents the optional [sequenced](#) list of one or more [constraint](#)
- 1276 [clauses](#) on the parameter definition.
- 1277 • **entry\_description**: represents the optional [description](#) of the entry schema.
- 1278 • **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.

### 1281 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.

### 1288 3.5.12.4 Example

1289 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 ]
```

1290 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 ] }
```

1291

### 1292 3.5.13 Operation definition

1293 An operation definition defines a named function or procedure that can be bound to an implementation  
1294 artifact (e.g., a script).

#### 1295 3.5.13.1 Keynames

1296 The following is the list of recognized keynames for a TOSCA operation definition:

Keyname	Required	Type	Description
description	no	<a href="#">description</a>	The optional description string for the associated named operation.
implementation	no	<a href="#">string</a>	The optional implementation artifact name (e.g., a script file name within a TOSCA CSAR file).
inputs	no	list of <a href="#">property definitions</a>	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.

Keyname	Required	Type	Description
	no	list of <a href="#">property assignments</a>	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.

1297 The following is the list of recognized keynames to be used with the **implementation** keyname within a  
1298 TOSCA operation definition:

Keyname	Required	Type	Description
primary	no	<a href="#">string</a>	The optional implementation artifact name (i.e., the primary script file name within a TOSCA CSAR file).
dependencies	no	list of <a href="#">string</a>	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).

### 1299 3.5.13.2 Grammar

1300 Operation definitions have the following grammars:

#### 1301 3.5.13.2.1 Short notation

1302 The following single-line grammar may be used when only an operation's implementation artifact is  
1303 needed:

```
<operation name>: <implementation artifact name>
```

#### 1304 3.5.13.2.2 Extended notation for use in Type definitions

1305 The following multi-line grammar may be used in Node or Relationship Type definitions when additional  
1306 information about the operation is needed:

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

#### 1307 3.5.13.2.3 Extended notation for use in Template definitions

1308 The following multi-line grammar may be used in Node or Relationship Template definitions when there  
1309 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>
```

1310 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1311 • **operation\_name**: represents the required symbolic name of the operation as a [string](#).
- 1312 • **operation\_description**: represents the optional [description](#) string for the corresponding
- 1313 **operation\_name**.
- 1314 • **implementation\_artifact\_name**: represents the optional name ([string](#)) of an implementation
- 1315 artifact definition (defined elsewhere), or the direct name of an implementation artifact's relative
- 1316 filename (e.g., a service template-relative, path-inclusive filename or absolute file location using a
- 1317 URL).
- 1318 • **property\_definitions**: represents the optional list of [property definitions](#) which the TOSCA
- 1319 orchestrator would make available (i.e., or pass) to the corresponding implementation artifact
- 1320 during its execution.
- 1321 • **property\_assignments**: represents the optional list of [property assignments](#) for passing
- 1322 parameters to Node or Relationship Template operations providing values for properties defined
- 1323 in their respective type definitions.
- 1324 • **list\_of\_dependent\_artifact\_names**: represents the optional ordered list of one or more
- 1325 dependent or secondary implementation artifact names (as strings) which are referenced by the
- 1326 primary implementation artifact. TOSCA orchestrators will copy these files to the same location
- 1327 as the primary artifact on the target node so as to make them accessible to the primary
- 1328 implementation artifact when it is executed.

### 1329 3.5.13.3 Additional requirements

- 1330 • The default sub-classing behavior for implementations of operations SHALL be override. That is,
- 1331 implementation artifacts assigned in subclasses override any defined in its parent class.
- 1332 • Template authors MAY provide property assignments on operation inputs on templates that do
- 1333 not necessarily have a property definition defined in its corresponding type.
- 1334 • Implementation artifact file names (e.g., script filenames) may include file directory path names
- 1335 that are relative to the TOSCA service template file itself when packaged within a TOSCA Cloud
- 1336 Service ARchive (CSAR) file.

### 1337 3.5.13.4 Examples

#### 1338 3.5.13.4.1 Single-line implementation example

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

#### 1339 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

```

1340 **3.5.14 Interface definition**

1341 An interface definition defines a named interface that can be associated with a Node or Relationship Type

1342 **3.5.14.1 Keynames**

1343 The following is the list of recognized keynames for a TOSCA interface definition:

Keyname	Required	Type	Description
inputs	no	list of <a href="#">property definitions</a>	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 <a href="#">property assignments</a>	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.

1344 **3.5.14.2 Grammar**

1345 Interface definitions have the following grammar:

1346 **3.5.14.2.1 Extended notation for use in Type definitions**

1347 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>

```

1348 **3.5.14.2.2 Extended notation for use in Template definitions**

1349 The following multi-line grammar may be used in Node or Relationship Template definitions:

```

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

```

1350 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1351 • **interface\_definition\_name:** represents the required symbolic name of the interface as a  
1352 [string](#).
- 1353 • **interface\_type\_name:** represents the required name of the Interface Type for the interface  
1354 **definition**.
- 1355 • **property\_definitions:** represents the optional list of [property definitions](#) (i.e., parameters)  
1356 which the TOSCA orchestrator would make available (i.e., or pass) to all defined operations.  
1357 - *This means these properties and their values would be accessible to the implementation*  
1358 *artifacts (e.g., scripts) associated to each operation during their execution.*
- 1359 • **property\_assignments:** represents the optional list of [property assignments](#) for passing  
1360 parameters to Node or Relationship Template operations providing values for properties defined  
1361 in their respective type definitions.
- 1362 • **operation\_definitions:** represents the required name of one or more [operation definitions](#).

### 1363 3.5.15 Event Filter definition

1364 An event filter definition defines criteria for selection of an attribute, for the purpose of monitoring it, within  
1365 a TOSCA entity, or one its capabilities.

#### 1366 3.5.15.1 Keynames

1367 The following is the list of recognized keynames for a TOSCA event filter definition:

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

#### 1368 3.5.15.2 Grammar

1369 Event filter definitions have following grammar:

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

1370 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1371 • **node\_type\_name:** represents the required name of the node type that would be used to select  
1372 (filter) the node that contains the attribute to monitor or contains the requirement that references  
1373 another node that contains the attribute to monitor.
- 1374 • **node\_template\_name:** represents the required name of the node template that would be used to  
1375 select (filter) the node that contains the attribute to monitor or contains the requirement that  
1376 references another node that contains the attribute to monitor.

- 1377 • **requirement\_name**: represents the optional name of the requirement that would be used to  
1378 select (filter) a referenced node that contains the attribute to monitor.
- 1379 • **capability\_name**: represents the optional name of a capability that would be used to select  
1380 (filter) the attribute to monitor.

### 1381 3.5.16 Trigger definition

1382 A trigger definition defines the event, condition and action that is used to “trigger” a policy it is associated  
1383 with.

#### 1384 3.5.16.1 Keynames

1385 The following is the list of recognized keynames for a TOSCA trigger definition:

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

#### 1386 3.5.16.2 Grammar

1387 Trigger definitions have the following grammars:

```

<trigger_name>:
  description: <trigger_description>
  # TBD: need to separate “simple” and “full” grammar for event type name
  event: <event_type_name>
    type: <event_type_name>
  schedule: <time_interval_for_trigger>

```

```

target_filter:
  <event_filter_definition>
condition: <attribute_constraint_clause>
  constraint: <constraint_clause>
  period: <scalar-unit.time> # e.g., 60 sec
  evaluations: <integer> # e.g., 1
  method: <string> # e.g., average
action:
  <operation_definition>

```

1388 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1389 • **trigger\_name**: represents the required symbolic name of the trigger as a [string](#).
- 1390 • **trigger\_description**: represents the optional [description](#) string for the corresponding
- 1391 **trigger\_name**.
- 1392 • **event\_type\_name**: represents the required name of the TOSCA Event Type that would be
- 1393 monitored on the identified resource (node).
- 1394 • **time\_interval\_for\_trigger**: represents the optional time interval that the trigger is valid
- 1395 **for**.
- 1396 • **event\_filter\_definition**: represents the optional filter to use to locate the resource (node)
- 1397 or capability attribute to monitor.
- 1398 • **attribute\_constraint\_clause**: represents the optional attribute constraint that would be
- 1399 used to test for a specific condition on the monitored resource.
- 1400 • **operation\_definition**: represents the required action to take if the event and (optionally)
- 1401 condition are met.

### 1402 3.5.17 Workflow activity definition

1403 A workflow activity defines an operation to be performed in a TOSCA workflow. Activities allows to:

- 1404
- 1405 • Delegate the workflow for a node expected to be provided by the orchestrator
- 1406 • Set the state of a node
- 1407 • Call an operation defined on a TOSCA interface of a node, relationship or group
- 1408 • Inline another workflow defined in the topology (to allow reusability)

#### 1409 3.5.17.1 Keynames

1410 The following is the list of recognized keynames for a TOSCA workflow activity definition. Note that while

1411 each of the key is not required, one and only one of them is required (mutually exclusive).

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

Keyname	Required	Type	Description
inline	no	string	The name of a workflow to be inlined.

1412 **3.5.17.2 Grammar**

1413 Workflow activity definitions have one of the following grammars:

1414 **3.5.17.2.1 Delegate activity**

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

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

- 1416 • **delegate\_workflow\_name**: represents the name of the workflow of the node
- 1417 provided by the TOSCA orchestrator.

1418 **3.5.17.2.2 Set state activity**

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

1419 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1420 • **new\_node\_state**: represents the state that will be affected to the node once
- 1421 the activity is performed.

1422 **3.5.17.2.3 Call operation activity:**

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

1423 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1424 • **interface\_name**: represents the name of the interface in which the operation to
- 1425 be called is defined.
- 1426 • **operation\_name**: represents the name of the operation of the interface that
- 1427 will be called during the workflow execution.

1428 **3.5.17.2.4 Inline activity**

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

1429 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1430 • **workflow\_name**: represents the name of the workflow to inline.

1431 **3.5.17.3 Additional Requirements**

- 1432 • Keynames are mutually exclusive, i.e. an activity MUST define only one of delegate, set\_state,
- 1433 call\_operation or inline keyname.

1434 **3.5.17.4 Example**

1435 following represents a list of workflow activity definitions:

```
- delegate: deploy
```

```
- set_state: started
- call_operation: toska.interfaces.node.lifecycle.Standard.start
- inline: my_workflow
```

1436

### 1437 3.5.18 Assertion definition

1438 A workflow assertion is used to specify a single condition on a workflow filter definition. The assertion  
1439 allows to assert the value of an attribute based on TOSCA constraints.

#### 1440 3.5.18.1 Keynames

1441 The TOSCA workflow assertion definition has no keynames.

#### 1442 3.5.18.2 Grammar

1443 Workflow assertion definitions have the following grammar:

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

1444 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1445 • **attribute\_name**: represents the name of an attribute defined on the assertion context entity  
1446 (node instance, relationship instance, group instance) and from which value will be evaluated  
1447 against the defined constraint clauses.
- 1448 • **list\_of\_constraint\_clauses**: represents the list of constraint clauses that will be used to validate  
1449 the attribute assertion.

#### 1450 3.5.18.3 Example

1451 Following represents a workflow assertion with a single equals constraint:

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

1452 Following represents a workflow assertion with multiple constraints:

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

### 1453 3.5.19 Condition clause definition

1454 A workflow condition clause definition is used to specify a condition that can be used within a workflow  
1455 precondition or workflow filter.

#### 1456 3.5.19.1 Keynames

1457 The following is the list of recognized keynames for a TOSCA workflow condition definition:

Keyname	Required	Type	Description
and	no	list of <a href="#">condition clause definition</a>	An <b>and</b> clause allows to define sub-filter clause definitions that must all be evaluated truly so the and clause is considered as true.

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

### 1458 3.5.19.2 Grammar

1459 Workflow assertion definitions have the following grammars:

#### 1460 3.5.19.2.1 And clause

```
and: <list_of_condition_clause_definition>
```

1461 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1462 • **list\_of\_condition\_clause\_definition:** represents the list of condition clauses. All  
1463 condition clauses MUST be asserted to true so that the and clause is asserted to true.

#### 1464 3.5.19.2.2 Or clause

```
or: <list_of_condition_clause_definition>
```

1465 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1466 • **list\_of\_condition\_clause\_definition:** represents the list of condition clauses. One of the  
1467 condition clause have to be asserted to true so that the or clause is asserted to true.

#### 1468 3.5.19.2.3 Assert clause

```
assert: <list_of_assertion_definition>
```

1469 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1470 • **list\_of\_assertion\_definition:** represents the list of assertions. All assertions MUST be  
1471 asserted to true so that the assert clause is asserted to true.

### 1472 3.5.19.3 Additional Requirement

- 1473 • Keynames are mutually exclusive, i.e. a filter definition can define only one of *and*, *or*, or *assert*  
1474 keyname.

### 1475 3.5.19.4 Notes

- 1476 • The TOSCA processor SHOULD perform assertion in the order of the list for every defined  
1477 condition clause or assertion definition.

### 1478 3.5.19.5 Example

1479 Following represents a workflow condition clause with a single equals constraint:

```

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

```

1480 Following represents a workflow condition clause with a single equals constraints on two different  
 1481 attributes:

```

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

```

1482 Following represents a workflow condition clause with a or constraint on two different assertions:

```

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

```

1483 Following represents multiple levels of condition clauses to build the following logic: one\_attribute equal  
 1484 one\_value AND (my\_attribute equal my\_value OR my\_other\_attribute equal my\_other\_value):

```

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

```

### 1485 3.5.20 Workflow precondition definition

1486 A workflow condition can be used as a filter or precondition to check if a workflow can be processed or  
 1487 not based on the state of the instances of a TOSCA topology deployment. When not met, the workflow  
 1488 will not be triggered.

#### 1489 3.5.20.1 Keynames

1490 The following is the list of recognized keynames for a TOSCA workflow condition definition:

Keyname	Required	Type	Description
target	yes	string	The target of the precondition (this can be a node template name, a group name)

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

1491 **3.5.20.2 Grammar**

1492 Workflow precondition definitions have the following grammars:

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

1493 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1494 • **target\_name:** represents the name of a node template or group in the topology.
- 1495 • **target\_requirement\_name:** represents the name of a requirement of the node template (in case
- 1496 target\_name refers to a node template.
- 1497 • **list\_of\_condition\_clause\_definition:** represents the list of condition clauses
- 1498 to be evaluated. The value of the resulting condition is evaluated as an AND
- 1499 clause between the different elements.

1500 **3.5.21 Workflow step definition**

1501 A workflow step allows to define one or multiple sequenced activities in a workflow and how they are

1502 connected to other steps in the workflow. They are the building blocks of a declarative workflow.

1503 **3.5.21.1 Keynames**

1504 The following is the list of recognized keynames for a TOSCA workflow step definition:

Keyname	Required	Type	Description
target	yes	string	The target of the step (this can be a node template name, a group name)
target_relationship	no	string	The optional name of a requirement of the target in case the step refers to a relationship rather than a node or group. Note that this is applicable only if the target is a node.

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

1505 **3.5.21.2 Grammar**

1506 Workflow step definitions have the following grammars:

```

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

```

1507 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1508 • **target\_name**: represents the name of a node template or group in the topology.
- 1509 • **target\_requirement\_name**: represents the name of a requirement of the node template (in case
- 1510 target\_name refers to a node template.
- 1511 • **operation\_host**: the node on which the operation should be executed
- 1512 • **<list\_of\_condition\_clause\_definition>**: represents a list of condition clause definition.

- 1513 • `list_of_activity_definition`: represents a list of activity definition
- 1514 • `target_step_name`: represents the name of another step of the workflow.

## 1515 3.6 Type-specific definitions

### 1516 3.6.1 Entity Type Schema

1517 An Entity Type is the common, base, polymorphic schema type which is extended by TOSCA base entity  
 1518 type schemas (e.g., Node Type, Relationship Type, Artifact Type, etc.) and serves to define once all the  
 1519 commonly shared keynames and their types. This is a “meta” type which is abstract and not directly  
 1520 instantiatable.

#### 1521 3.6.1.1 Keynames

1522 The following is the list of recognized keynames for a TOSCA Entity Type definition:

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

#### 1523 3.6.1.2 Grammar

1524 Entity Types have following grammar:

```

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

1525 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1526 • `version_number`: represents the optional TOSCA `version` number for the entity.
- 1527 • `entity_description`: represents the optional `description` string for the entity.
- 1528 • `metadata_map`: represents the optional map of string.

#### 1529 3.6.1.3 Additional Requirements

- 1530 • The TOSCA Entity Type SHALL be the common base type used to derive all other top-level base  
 1531 TOSCA Types.
- 1532 • The TOSCA Entity Type SHALL NOT be used to derive or create new base types apart from  
 1533 those defined in this specification or a profile of this specification.

1534 **3.6.2 Capability definition**

1535 A capability definition defines a named, typed set of data that can be associated with Node Type or Node  
1536 Template to describe a transparent capability or feature of the software component the node describes.

1537 **3.6.2.1 Keynames**

1538 The following is the list of recognized keynames for a TOSCA capability definition:

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

1539 **3.6.2.2 Grammar**

1540 Capability definitions have one of the following grammars:

1541 **3.6.2.2.1 Short notation**

1542 The following grammar may be used when only a list of capability definition names needs to be declared:

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

1543 **3.6.2.2.2 Extended notation**

1544 The following multi-line grammar may be used when additional information on the capability definition is  
1545 needed:

```
<capability_definition_name>:  
  type: <capability_type>  
  description: <capability_description>  
  properties:  
    <property_definitions>  
  attributes:
```

```
<attribute_definitions>
valid_source_types: [ <node_type_names> ]
```

1546 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1547 • **capability\_definition\_name**: represents the symbolic name of the capability as a [string](#).
- 1548 • **capability\_type**: represents the required name of a [capability type](#) the capability definition is
- 1549 based upon.
- 1550 • **capability\_description**: represents the optional [description](#) of the capability definition.
- 1551 • **property\_definitions**: represents the optional list of [property definitions](#) for the capability
- 1552 definition.
- 1553 • **attribute\_definitions**: represents the optional list of [attribute definitions](#) for the capability
- 1554 definition.
- 1555 • **node\_type\_names**: represents the optional list of one or more names of [Node Types](#) that the
- 1556 Capability definition supports as valid sources for a successful relationship to be established to
- 1557 itself.

### 1558 3.6.2.3 Examples

1559 The following examples show capability definitions in both simple and full forms:

#### 1560 3.6.2.3.1 Simple notation example

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

#### 1561 3.6.2.3.2 Full notation example

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

### 1562 3.6.2.4 Additional requirements

- 1563 • Any Node Type (names) provides as values for the **valid\_source\_types** keyname SHALL be
- 1564 type-compatible (i.e., derived from the same parent Node Type) with any Node Types defined
- 1565 using the same keyname in the parent Capability Type.
- 1566 • Capability symbolic names SHALL be unique; it is an error if a capability name is found to occur
- 1567 more than once.

### 1568 3.6.2.5 Notes

- 1569 • The Capability Type, in this example **MyCapabilityTypeName**, would be defined
- 1570 elsewhere and have an integer property named **limit**.
- 1571 • This definition directly maps to the **CapabilitiesDefinition** of the Node Type entity as defined
- 1572 in the [TOSCA v1.0 specification](#).

### 1573 3.6.3 Requirement definition

1574 The Requirement definition describes a named requirement (dependencies) of a TOSCA Node Type or  
1575 Node template which needs to be fulfilled by a matching Capability definition declared by another TOSCA  
1576 modelable entity. The requirement definition may itself include the specific name of the fulfilling entity  
1577 (explicitly) or provide an abstract type, along with additional filtering characteristics, that a TOSCA  
1578 orchestrator can use to fulfill the capability at runtime (implicitly).

#### 1579 3.6.3.1 Keynames

1580 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 <a href="#">Capability Type</a> that can fulfill the requirement.
node	no	string	N/A	The optional reserved keyname used to provide the name of a valid <a href="#">Node Type</a> that contains the capability definition that can be used to fulfill the requirement.
relationship	no	string	N/A	The optional reserved keyname used to provide the name of a valid <a href="#">Relationship Type</a> 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 <b>UNBOUNDED</b> is also supported to represent any positive integer.

#### 1581 3.6.3.1.1 Additional Keynames for multi-line relationship grammar

1582 The Requirement definition contains the Relationship Type information needed by TOSCA Orchestrators  
1583 to construct relationships to other TOSCA nodes with matching capabilities; however, it is sometimes  
1584 recognized that additional properties may need to be passed to the relationship (perhaps for  
1585 configuration). In these cases, additional grammar is provided so that the Node Type may declare  
1586 additional Property definitions to be used as inputs to the Relationship Type's declared interfaces (or  
1587 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 <b>relationship</b> 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.

#### 1588 3.6.3.2 Grammar

1589 Requirement definitions have one of the following grammars:

##### 1590 3.6.3.2.1 Simple grammar (Capability Type only)

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

### 1591 3.6.3.2.2 Extended grammar (with Node and Relationship Types)

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

### 1592 3.6.3.2.3 Extended grammar for declaring Property Definitions on the 1593 relationship's Interfaces

1594 The following additional multi-line grammar is provided for the relationship keyname in order to declare  
1595 new Property definitions for inputs of known Interface definitions of the declared Relationship Type.

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

1596 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1597 • **requirement\_name**: represents the required symbolic name of the requirement definition as a  
1598 **string**.
- 1599 • **capability\_type\_name**: represents the required name of a Capability type that can be used to  
1600 fulfill the requirement.
- 1601 • **node\_type\_name**: represents the optional name of a TOSCA Node Type that contains the  
1602 Capability Type definition the requirement can be fulfilled by.
- 1603 • **relationship\_type\_name**: represents the optional name of a **Relationship Type** to be used to  
1604 construct a relationship between this requirement definition (i.e., in the source node) to a  
1605 matching capability definition (in a target node).
- 1606 • **min\_occurrences**, **max\_occurrences**: represents the optional minimum and maximum  
1607 occurrences of the requirement (i.e., its cardinality).
- 1608 • **interface\_definitions**: represents one or more already declared interface definitions in the  
1609 Relationship Type (as declared on the **type** keyname) allowing for the declaration of new  
1610 Property definition for these interfaces or for specific Operation definitions of these interfaces.

### 1611 3.6.3.3 Additional Requirements

- 1612 • Requirement symbolic names SHALL be unique; it is an error if a requirement name is found to  
1613 occur more than once.
- 1614 • If the **occurrences** keyname is not present, then the occurrence of the requirement **SHALL** be  
1615 one and only one; that is a default declaration as follows would be assumed:  
1616     o occurrences: [1,1]

### 1617 3.6.3.4 Notes

- 1618 • This element directly maps to the **RequirementsDefinition** of the Node Type entity as defined  
1619 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.

### 3.6.3.5 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.

#### 3.6.3.5.1 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.

## 3.6.4 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.

### 3.6.4.1 Keynames

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

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

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

### 3.6.4.2 Grammar

Artifact Types have following grammar:

```

<artifact_type_name>:
  derived_from: <parent_artifact_type_name>
  version: <version_number>
  metadata:
    <map of string>
  description: <artifact_description>
  mime_type: <mime_type_string>

```

```
file_ext: [ <file_extensions> ]
properties:
  <property_definitions>
```

1646 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1647 • **artifact\_type\_name**: represents the name of the Artifact Type being declared as a [string](#).
- 1648 • **parent\_artifact\_type\_name**: represents the [name](#) of the [Artifact Type](#) this Artifact Type definition derives from (i.e., its “parent” type).
- 1649 • **version\_number**: represents the optional TOSCA [version](#) number for the Artifact Type.
- 1650 • **artifact\_description**: represents the optional [description](#) string for the Artifact Type.
- 1651 • **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](#).
- 1652 • **file\_extensions**: represents the optional list of one or more recognized file extensions for this type of artifact type as [strings](#).
- 1653 • **property\_definitions**: represents the optional list of [property definitions](#) for the artifact type.

### 1657 3.6.4.3 Examples

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

### 1658 3.6.4.4 Notes

- 1659 • The ‘mime\_type’ keyname is meant to have values that are Apache mime types such as those defined here: <http://svn.apache.org/repos/asf/httpd/httpd/trunk/docs/conf/mime.types>

## 1661 3.6.5 Interface Type

1662 An Interface Type is a reusable entity that describes a set of operations that can be used to interact with  
1663 or manage a node or relationship in a TOSCA topology.

### 1664 3.6.5.1 Keynames

1665 The Interface Type is a TOSCA Entity and has the common keynames listed in section 3.6.1 TOSCA  
1666 Entity Schema.

1667 In addition, the Interface Type has the following recognized keynames:

Keyname	Required	Type	Description
inputs	no	list of <a href="#">property definitions</a>	The optional list of input parameter definitions.

### 1668 3.6.5.2 Grammar

1669 Interface Types have following grammar:

```

<interface_type_name>:
  derived_from: <parent_interface_type_name>
  version: <version_number>
  metadata:
    <map of string>
  description: <interface_description>
  inputs:
    <property_definitions>
    <operation_definitions>

```

1670 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1671 • **interface\_type\_name**: represents the required name of the interface as a [string](#).
- 1672 • **parent\_interface\_type\_name**: represents the name of the [Interface Type](#) this Interface Type
- 1673 definition derives from (i.e., its “parent” type).
- 1674 • **version\_number**: represents the optional TOSCA [version](#) number for the Interface Type.
- 1675 • **interface\_description**: represents the optional [description](#) string for the Interface Type.
- 1676 • **property\_definitions**: represents the optional list of [property definitions](#) (i.e., parameters)
- 1677 which the TOSCA orchestrator would make available (i.e., or pass) to all implementation artifacts
- 1678 for operations declared on the interface during their execution.
- 1679 • **operation\_definitions**: represents the required list of one or more [operation definitions](#).

### 1680 3.6.5.3 Example

1681 The following example shows a custom interface used to define multiple configure operations.

```

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

```

### 1682 3.6.5.4 Additional Requirements

- 1683 • Interface Types **MUST NOT** include any implementations for defined operations; that is, the
- 1684 implementation keyname is invalid.
- 1685 • The **inputs** keyname is reserved and **SHALL NOT** be used for an operation name.

### 1686 3.6.6 Data Type

1687 A Data Type definition defines the schema for new named datatypes in TOSCA.

1688 **3.6.6.1 Keynames**

1689 The Data Type is a TOSCA Entity and has the common keynames listed in section 3.6.1 TOSCA Entity  
1690 Schema.

1691 In addition, the Data Type has the following recognized keynames:

Keyname	Required	Type	Description
constraints	no	list of <a href="#">constraint clauses</a>	The optional list of <i>sequenced</i> constraint clauses for the Data Type.
properties	no	list of <a href="#">property definitions</a>	The optional list property definitions that comprise the schema for a complex Data Type in TOSCA.

1692 **3.6.6.2 Grammar**

1693 Data Types have the following grammar:

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

1694 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1695 • **data\_type\_name**: represents the required symbolic name of the Data Type as a [string](#).
- 1696 • **version\_number**: represents the optional TOSCA [version](#) number for the Data Type.
- 1697 • **datatype\_description**: represents the optional [description](#) for the Data Type.
- 1698 • **existing\_type\_name**: represents the optional name of a valid TOSCA type this new Data  
1699 Type would derive from.
- 1700 • **type\_constraints**: represents the optional *sequenced* list of one or more type-compatible  
1701 [constraint clauses](#) that restrict the Data Type.
- 1702 • **property\_definitions**: represents the optional list of one or more [property definitions](#) that  
1703 provide the schema for the Data Type.

1704 **3.6.6.3 Additional Requirements**

- 1705 • A valid datatype definition **MUST** have either a valid **derived\_from** declaration or at least one  
1706 valid property definition.
- 1707 • Any **constraint** clauses **SHALL** be type-compatible with the type declared by the  
1708 **derived\_from** keyname.
- 1709 • If a **properties** keyname is provided, it **SHALL** contain one or more valid property definitions.

1710 **3.6.6.4 Examples**

1711 The following example represents a Data Type definition based upon an existing string type:

### 1712 3.6.6.4.1 Defining a complex datatype

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

### 1713 3.6.6.4.2 Defining a datatype derived from an existing datatype

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

## 1714 3.6.7 Capability Type

1715 A Capability Type is a reusable entity that describes a kind of capability that a Node Type can declare to  
1716 expose. Requirements (implicit or explicit) that are declared as part of one node can be matched to (i.e.,  
1717 fulfilled by) the Capabilities declared by another node.

### 1718 3.6.7.1 Keynames

1719 The Capability Type is a TOSCA Entity and has the common keynames listed in section 3.6.1 TOSCA  
1720 Entity Schema.

1721 In addition, the Capability Type has the following recognized keynames:

Keyname	Required	Type	Description
properties	no	list of <a href="#">property definitions</a>	An optional list of property definitions for the Capability Type.
attributes	no	list of <a href="#">attribute definitions</a>	An optional list of attribute definitions for the Capability Type.
valid_source_types	no	<a href="#">string[]</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.

### 1722 3.6.7.2 Grammar

1723 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> ]
```

1724 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1725 • **capability\_type\_name**: represents the required name of the Capability Type being declared as  
1726 a **string**.
- 1727 • **parent\_capability\_type\_name**: represents the name of the **Capability Type** this Capability  
1728 Type definition derives from (i.e., its “parent” type).
- 1729 • **version\_number**: represents the optional TOSCA **version** number for the Capability Type.
- 1730 • **capability\_description**: represents the optional **description** string for the corresponding  
1731 **capability\_type\_name**.
- 1732 • **property\_definitions**: represents an optional list of **property definitions** that the Capability  
1733 type exports.
- 1734 • **attribute\_definitions**: represents the optional list of **attribute definitions** for the Capability  
1735 Type.
- 1736 • **node\_type\_names**: represents the optional list of one or more names of **Node Types** that the  
1737 Capability Type supports as valid sources for a successful relationship to be established to itself.

### 1738 3.6.7.3 Example

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

### 1739 3.6.8 Requirement Type

1740 A Requirement Type is a reusable entity that describes a kind of requirement that a Node Type can  
1741 declare to expose. The TOSCA Simple Profile seeks to simplify the need for declaring specific  
1742 Requirement Types from nodes and instead rely upon nodes declaring their features sets using TOSCA  
1743 Capability Types along with a named Feature notation.

1744 Currently, there are no use cases in this TOSCA Simple Profile in YAML specification that utilize an  
 1745 independently defined Requirement Type. This is a desired effect as part of the simplification of the  
 1746 TOSCA v1.0 specification.

### 1747 3.6.9 Node Type

1748 A Node Type is a reusable entity that defines the type of one or more Node Templates. As such, a Node  
 1749 Type defines the structure of observable properties via a *Properties Definition, the Requirements and*  
 1750 *Capabilities of the node as well as its supported interfaces.*

#### 1751 3.6.9.1 Keynames

1752 The Node Type is a TOSCA Entity and has the common keynames listed in section 3.6.1 TOSCA Entity  
 1753 Schema.

1754 In addition, the Node Type has the following recognized keynames:

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

#### 1755 3.6.9.2 Grammar

1756 Node Types have following grammar:

```

<node_type_name>:
  derived_from: <parent_node_type_name>
  version: <version_number>
  metadata:
    <map of string>
  description: <node_type_description>
  attributes:
    <attribute_definitions>
  properties:
    <property_definitions>
  requirements:
    - <requirement_definitions>
  capabilities:
  
```

```

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

```

1757 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1758 • **node\_type\_name**: represents the required symbolic name of the Node Type being declared.
- 1759 • **parent\_node\_type\_name**: represents the name (*string*) of the **Node Type** this Node Type
- 1760 definition derives from (i.e., its “parent” type).
- 1761 • **version\_number**: represents the optional TOSCA **version** number for the Node Type.
- 1762 • **node\_type\_description**: represents the optional **description** string for the corresponding
- 1763 **node\_type\_name**.
- 1764 • **property\_definitions**: represents the optional list of **property definitions** for the Node Type.
- 1765 • **attribute\_definitions**: represents the optional list of **attribute definitions** for the Node Type.
- 1766 • **requirement\_definitions**: represents the optional *sequenced* list of **requirement definitions** for
- 1767 the Node Type.
- 1768 • **capability\_definitions**: represents the optional list of **capability definitions** for the Node
- 1769 Type.
- 1770 • **interface\_definitions**: represents the optional list of one or more **interface definitions**
- 1771 supported by the Node Type.
- 1772 • **artifact\_definitions**: represents the optional list of **artifact definitions** for the Node Type.

### 1773 3.6.9.3 Additional Requirements

- 1774 • Requirements are intentionally expressed as a sequenced list of TOSCA **Requirement definitions**
- 1775 which **SHOULD** be resolved (processed) in sequence order by TOSCA Orchestrators. .

### 1776 3.6.9.4 Best Practices

- 1777 • It is recommended that all Node Types **SHOULD** derive directly (as a parent) or indirectly (as an
- 1778 ancestor) of the TOSCA Root Node Type (i.e., **tosca.nodes.Root**) to promote compatibility and
- 1779 portability. However, it is permitted to author Node Types that do not do so.
- 1780 • TOSCA Orchestrators, having a full view of the complete application topology template and its
- 1781 resultant dependency graph of nodes and relationships, **MAY** prioritize how they instantiate the nodes
- 1782 and relationships for the application (perhaps in parallel where possible) to achieve the greatest
- 1783 efficiency

### 1784 3.6.9.5 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

```

1785 **3.6.10 Relationship Type**

1786 A Relationship Type is a reusable entity that defines the type of one or more relationships between Node  
 1787 Types or Node Templates.

1788 **3.6.10.1 Keynames**

1789 The Relationship Type is a TOSCA Entity and has the common keynames listed in section 3.6.1 TOSCA  
 1790 Entity Schema.

1791 In addition, the Relationship Type has the following recognized keynames:

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

1792 **3.6.10.2 Grammar**

1793 Relationship Types have following grammar:

```

<relationship type name>:
  derived_from: <parent relationship type name>
  version: <version number>
  metadata:
    <map of string>
  description: <relationship description>

```

```
properties:
  <property_definitions>
attributes:
  <attribute_definitions>
interfaces:
  <interface_definitions>
valid_target_types: [ <capability_type_names> ]
```

1794 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1795 • **relationship\_type\_name**: represents the required symbolic name of the Relationship Type  
1796 being declared as a [string](#).
- 1797 • **parent\_relationship\_type\_name**: represents the name ([string](#)) of the [Relationship Type](#) this  
1798 Relationship Type definition derives from (i.e., its “parent” type).
- 1799 • **relationship\_description**: represents the optional [description](#) string for the corresponding  
1800 **relationship\_type\_name**.
- 1801 • **version\_number**: represents the optional TOSCA [version](#) number for the Relationship Type.
- 1802 • **property\_definitions**: represents the optional list of [property definitions](#) for the Relationship  
1803 Type.
- 1804 • **attribute\_definitions**: represents the optional list of [attribute definitions](#) for the Relationship  
1805 Type.
- 1806 • **interface\_definitions**: represents the optional list of one or more names of valid [interface](#)  
1807 [definitions](#) supported by the Relationship Type.
- 1808 • **capability\_type\_names**: represents one or more names of valid target types for the  
1809 relationship (i.e., [Capability Types](#)).

### 1810 3.6.10.3 Best Practices

- 1811 • For TOSCA application portability, it is recommended that designers use the normative  
1812 Relationship types defined in this specification where possible and derive from them for  
1813 customization purposes.
- 1814 • The TOSCA Root Relationship Type (**tosca.relationships.Root**) SHOULD be used to derive  
1815 new types where possible when defining new relationships types. This assures that its normative  
1816 configuration interface (**tosca.interfaces.relationship.Configure**) can be used in a  
1817 deterministic way by TOSCA orchestrators.

### 1818 3.6.10.4 Examples

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

### 1819 3.6.11 Group Type

1820 A Group Type defines logical grouping types for nodes, typically for different management purposes.  
1821 Groups can effectively be viewed as logical nodes that are not part of the physical deployment topology of  
1822 an application, yet can have capabilities and the ability to attach policies and interfaces that can be  
1823 applied (depending on the group type) to its member nodes.

1824

1825 Conceptually, group definitions allow the creation of logical “membership” relationships to nodes in a  
 1826 service template that are not a part of the application’s explicit requirement dependencies in the topology  
 1827 template (i.e. those required to actually get the application deployed and running). Instead, such logical  
 1828 membership allows for the introduction of things such as group management and uniform application of  
 1829 policies (i.e., requirements that are also not bound to the application itself) to the group’s members.

### 1830 3.6.11.1 Keynames

1831 The Group Type is a TOSCA Entity and has the common keynames listed in section 3.6.1 TOSCA Entity  
 1832 Schema.

1833 In addition, the Group Type has the following recognized keynames:

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

### 1834 3.6.11.2 Grammar

1835 Group Types have one the following grammars:

```

<group type name>:
  derived_from: <parent group type name>
  version: <version number>
  metadata:
    <map of string>
  description: <group description>
  properties:
    <property definitions>
  members: [ <list\_of\_valid\_member\_types> ]
  requirements:
    - <requirement definitions>
  capabilities:
  
```

```

<capability_definitions>
interfaces:
  <interface_definitions>

```

1836 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1837 • **group\_type\_name**: represents the required symbolic name of the Group Type being declared as
- 1838 a [string](#).
- 1839 • **parent\_group\_type\_name**: represents the name ([string](#)) of the [Group Type](#) this Group Type
- 1840 definition derives from (i.e., its “parent” type).
- 1841 • **version\_number**: represents the optional TOSCA [version](#) number for the Group Type.
- 1842 • **group\_description**: represents the optional description string for the corresponding
- 1843 **group\_type\_name**.
- 1844 • **property\_definitions**: represents the optional list of [property definitions](#) for the Group Type.
- 1845 • **list\_of\_valid\_member\_types**: represents the optional list of TOSCA types (e.g.,, Node,
- 1846 Capability or even other Group Types) that are valid member types for being added to (i.e.,
- 1847 members of) the Group Type.
- 1848 • **interface\_definitions**: represents the optional list of one or more [interface definitions](#)
- 1849 supported by the Group Type.

### 1850 3.6.11.3 Additional Requirements

- 1851 • Group definitions **SHOULD NOT** be used to define or redefine relationships (dependencies) for
- 1852 an application that can be expressed using normative TOSCA Relationships within a TOSCA
- 1853 topology template.
- 1854 • The list of values associated with the “members” keyname **MUST** only contain types that or
- 1855 homogenous (i.e., derive from the same type hierarchy).

### 1856 3.6.11.4 Example

1857 The following represents a Group Type definition:

```

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

```

### 1858 3.6.12 Policy Type

1859 A Policy Type defines a type of requirement that affects or governs an application or service’s topology at  
 1860 some stage of its lifecycle, but is not explicitly part of the topology itself (i.e., it does not prevent the  
 1861 application or service from being deployed or run if it did not exist).

#### 1862 3.6.12.1 Keynames

1863 The Policy Type is a TOSCA Entity and has the common keynames listed in section 3.6.1 TOSCA Entity  
 1864 Schema.

1865 In addition, the Policy Type has the following recognized keynames:

Keyname	Required	Type	Description
properties	no	list of <a href="#">property definitions</a>	An optional list of property definitions for the Policy Type.

Keyname	Required	Type	Description
targets	no	<a href="#">string</a> []	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).
triggers	no	list of <a href="#">trigger</a>	An optional list of policy triggers for the Policy Type.

1866 **3.6.12.2 Grammar**

1867 Policy Types have the following grammar:

```

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

```

1868 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1869 • **policy\_type\_name**: represents the required symbolic name of the Policy Type being declared
- 1870 as a [string](#).
- 1871 • **parent\_policy\_type\_name**: represents the name ([string](#)) of the Policy Type this Policy Type
- 1872 definition derives from (i.e., its "parent" type).
- 1873 • **version\_number**: represents the optional TOSCA [version](#) number for the Policy Type.
- 1874 • **policy\_description**: represents the optional description string for the corresponding
- 1875 **policy\_type\_name**.
- 1876 • **property\_definitions**: represents the optional list of [property definitions](#) for the Policy Type.
- 1877 • **list\_of\_valid\_target\_types**: represents the optional list of TOSCA types (i.e., Group or
- 1878 Node Types) that are valid targets for this Policy Type.
- 1879 • **list\_of\_trigger\_definitions**: represents the optional list of [trigger definitions](#) for the policy.

1880 **3.6.12.3 Example**

1881 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

```

## 1882 3.7 Template-specific definitions

1883 The definitions in this section provide reusable modeling element grammars that are specific to the Node  
1884 or Relationship templates.

### 1885 3.7.1 Capability assignment

1886 A capability assignment allows node template authors to assign values to properties and attributes for a  
1887 named capability definition that is part of a Node Template's type definition.

#### 1888 3.7.1.1 Keynames

1889 The following is the list of recognized keynames for a TOSCA capability assignment:

Keyname	Required	Type	Description
properties	no	list of <a href="#">property assignments</a>	An optional list of property definitions for the Capability definition.
attributes	no	list of <a href="#">attribute assignments</a>	An optional list of attribute definitions for the Capability definition.

#### 1890 3.7.1.2 Grammar

1891 Capability assignments have one of the following grammars:

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

1892 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1893 • **capability\_definition\_name**: represents the symbolic name of the capability as a [string](#).
- 1894 • **property\_assignments**: represents the optional list of [property assignments](#) for the capability  
1895 definition.
- 1896 • **attribute\_assignments**: represents the optional list of [attribute assignments](#) for the capability  
1897 definition.

#### 1898 3.7.1.3 Example

1899 The following example shows a capability assignment:

##### 1900 3.7.1.3.1 Notation example

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

1901 **3.7.2 Requirement assignment**

1902 A Requirement assignment allows template authors to provide either concrete names of TOSCA  
 1903 templates or provide abstract selection criteria for providers to use to find matching TOSCA templates  
 1904 that are used to fulfill a named requirement's declared TOSCA Node Type.

1905 **3.7.2.1 Keynames**

1906 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"> <li>• <b>Capability definition</b> within a <i>target</i> node template that can fulfill the requirement.</li> <li>• <b>Capability Type</b> that the provider will use to select a type-compatible <i>target</i> node template to fulfill the requirement at runtime.</li> </ul>
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"> <li>• <b>Node Template</b> name that can fulfill the target node requirement.</li> <li>• <b>Node Type</b> name that the provider will use to select a type-compatible node template to fulfill the requirement at runtime.</li> </ul>
relationship	no	string	The optional reserved keyname used to provide the name of either a: <ul style="list-style-type: none"> <li>• <b>Relationship Template</b> to use to relate the <i>source</i> node to the (capability in the) <i>target</i> node when fulfilling the requirement.</li> <li>• <b>Relationship Type</b> 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.</li> </ul>
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.

1907 The following is the list of recognized keynames for a TOSCA requirement assignment's **relationship**  
 1908 keyname which is used when Property assignments need to be provided to inputs of declared interfaces  
 1909 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 <b>relationship</b> keyname.
properties	no	list of interface definitions	The optional reserved keyname used to reference declared (named) interface definitions of the corresponding Relationship Type in order to provide Property assignments for these interfaces or operations of these interfaces.

1910 **3.7.2.2 Grammar**

1911 Named requirement assignments have one of the following grammars:

1912 **3.7.2.2.1 Short notation:**

1913 The following single-line grammar may be used if only a concrete Node Template for the target node  
 1914 needs to be declared in the requirement:

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

1915 This notation is only valid if the corresponding Requirement definition in the Node Template's parent  
1916 Node Type declares (at a minimum) a valid Capability Type which can be found in the declared target  
1917 Node Template. A valid capability definition always needs to be provided in the requirement declaration of  
1918 the *source* node to identify a specific capability definition in the *target* node the requirement will form a  
1919 TOSCA relationship with.

### 1920 3.7.2.2.2 Extended notation:

1921 The following grammar would be used if the requirement assignment needs to provide more information  
1922 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 ]
```

### 1923 3.7.2.2.3 Extended grammar with Property Assignments for the relationship's 1924 Interfaces

1925 The following additional multi-line grammar is provided for the relationship keyname in order to provide  
1926 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>
```

1927 Examples of uses for the extended requirement assignment grammar include:

- 1928 • The need to allow runtime selection of the target node based upon an abstract Node Type rather  
1929 than a concrete Node Template. This may include use of the `node_filter` keyname to provide  
1930 node and capability filtering information to find the “best match” of a concrete Node Template at  
1931 runtime.
- 1932 • The need to further clarify the concrete Relationship Template or abstract Relationship Type to  
1933 use when relating the source node's requirement to the target node's capability.
- 1934 • The need to further clarify the concrete capability (symbolic) name or abstract Capability Type in  
1935 the target node to form a relationship between.
- 1936 • The need to (further) constrain the occurrences of the requirement in the instance model.

1937 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1938 • **requirement\_name**: represents the symbolic name of a requirement assignment as a [string](#).

- 1939 • **node\_template\_name**: represents the optional name of a Node Template that contains the
- 1940 capability this requirement will be fulfilled by.
- 1941 • **relationship\_template\_name**: represents the optional name of a [Relationship Type](#) to be used
- 1942 when relating the requirement appears to the capability in the target node.
- 1943 • **capability\_symbolic\_name**: represents the optional ordered list of specific, required capability
- 1944 type or named capability definition within the target Node Type or Template.
- 1945 • **node\_type\_name**: represents the optional name of a TOSCA Node Type the associated named
- 1946 requirement can be fulfilled by. This must be a type that is compatible with the Node Type
- 1947 declared on the matching requirement (same symbolic name) the requirement's Node Template
- 1948 is based upon.
- 1949 • **relationship\_type\_name**: represents the optional name of a [Relationship Type](#) that is
- 1950 compatible with the Capability Type in the target node.
- 1951 • **property\_assignments**: represents the optional list of property value assignments for the
- 1952 declared relationship.
- 1953 • **interface\_assignments**: represents the optional list of interface definitions for the declared
- 1954 relationship used to provide property assignments on inputs of interfaces and operations.
- 1955 • **capability\_type\_name**: represents the optional name of a Capability Type definition within the
- 1956 target Node Type this requirement needs to form a relationship with.
- 1957 • **node\_filter\_definition**: represents the optional [node filter](#) TOSCA orchestrators would use
- 1958 to fulfill the requirement for selecting a target node. Note that this SHALL only be valid if the **node**
- 1959 keyname's value is a Node Type and is invalid if it is a Node Template.

### 1960 3.7.2.3 Examples

#### 1961 3.7.2.3.1 Example 1 – Abstract hosting requirement on a Node Type

1962 A web application node template named 'my\_application\_node\_template' of type **WebApplication**

1963 declares a requirement named 'host' that needs to be fulfilled by any node that derives from the node

1964 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
```

1965 In this case, the node template's type is **WebApplication** which already declares the Relationship Type

1966 **HostedOn** to use to relate to the target node and the Capability Type of **Container** to be the specific

1967 target of the requirement in the target node.

#### 1968 3.7.2.3.2 Example 2 - Requirement with Node Template and a custom Relationship

#### 1969 Type

1970 This example is similar to the previous example; however, the requirement named 'database' describes

1971 a requirement for a connection to a database endpoint (**Endpoint.Database**) Capability Type in a named

1972 node template (**my\_database**). However, the connection requires a custom Relationship Type

1973 (**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
```

1974 **3.7.2.3.3 Example 3 - Requirement for a Compute node with additional selection**  
 1975 **criteria (filter)**

1976 This example shows how to extend an abstract 'host' requirement for a **Compute** node with a filter  
 1977 definition that further constrains TOSCA orchestrators to include additional properties and capabilities  
 1978 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 ] }
```

1979 **3.7.3 Node Template**

1980 A Node Template specifies the occurrence of a manageable software component as part of an  
 1981 application's topology model which is defined in a TOSCA Service Template. A Node template is an  
 1982 instance of a specified Node Type and can provide customized properties, constraints or operations  
 1983 which override the defaults provided by its Node Type and its implementations.

1984 **3.7.3.1 Keynames**

1985 The following is the list of recognized keynames for a TOSCA Node Template definition:

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

### 1986 3.7.3.2 Grammar

```

<node template name>:
  type: <node type name>
  description: <node template description>
  directives: [<directives>]
  metadata:
    <map of string>
  properties:
    <property assignments>
  attributes:
    <attribute assignments>
  requirements:
    - <requirement assignments>
  capabilities:

```

```

    <capability_assignments>
  interfaces:
    <interface_definitions>
  artifacts:
    <artifact_definitions>
  node_filter:
    <node_filter_definition>
  copy: <source_node_template_name>

```

1987 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1988 • **node\_template\_name**: represents the required symbolic name of the Node Template being
- 1989 declared.
- 1990 • **node\_type\_name**: represents the name of the Node Type the Node Template is based upon.
- 1991 • **node\_template\_description**: represents the optional [description](#) string for Node Template.
- 1992 • **directives**: represents the optional list of processing instruction keywords (as strings) for use by
- 1993 tooling and orchestrators.
- 1994 • **property\_assignments**: represents the optional list of [property assignments](#) for the Node
- 1995 Template that provide values for properties defined in its declared Node Type.
- 1996 • **attribute\_assignments**: represents the optional list of [attribute assignments](#) for the Node
- 1997 Template that provide values for attributes defined in its declared Node Type.
- 1998 • **requirement\_assignments**: represents the optional *sequenced* list of [requirement assignments](#)
- 1999 for the Node Template that allow assignment of type-compatible capabilities, target nodes,
- 2000 relationships and target (node filters) for use when fulfilling the requirement at runtime.
- 2001 • **capability\_assignments**: represents the optional list of [capability assignments](#) for the Node
- 2002 Template that augment those provided by its declared Node Type.
- 2003 • **interface\_definitions**: represents the optional list of [interface definitions](#) for the Node
- 2004 Template that *augment* those provided by its declared Node Type.
- 2005 • **artifact\_definitions**: represents the optional list of [artifact definitions](#) for the Node Template
- 2006 that augment those provided by its declared Node Type.
- 2007 • **node\_filter\_definition**: represents the optional [node filter](#) TOSCA orchestrators would use
- 2008 for selecting a matching node template.
- 2009 • **source\_node\_template\_name**: represents the optional (symbolic) name of another node
- 2010 template to copy into (all keynames and values) and use as a basis for this node template.

### 2011 3.7.3.3 Additional requirements

- 2012 • The **node\_filter** keyword (and supporting grammar) **SHALL** only be valid if the Node Template
- 2013 has a **directive** keyname with the value of “**selectable**” set.
- 2014 • The source node template provided as a value on the **copy** keyname **MUST NOT** itself use the
- 2015 **copy** keyname (i.e., it must itself be a complete node template description and not copied from
- 2016 another node template).

### 2017 3.7.3.4 Example

```

node_templates:
  mysql:

```

```

type: toska.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

```

2018 **3.7.4 Relationship Template**

2019 A Relationship Template specifies the occurrence of a manageable relationship between node templates  
 2020 as part of an application's topology model that is defined in a TOSCA Service Template. A Relationship  
 2021 template is an instance of a specified Relationship Type and can provide customized properties,  
 2022 constraints or operations which override the defaults provided by its Relationship Type and its  
 2023 implementations.

2024 **3.7.4.1 Keynames**

2025 The following is the list of recognized keynames for a TOSCA Relationship Template definition:

Keyname	Required	Type	Description
type	yes	<a href="#">string</a>	The required name of the Relationship Type the Relationship Template is based upon.
description	no	<a href="#">description</a>	An optional description for the Relationship Template.
metadata	no	<a href="#">map of string</a>	Defines a section used to declare additional metadata information.
properties	no	list of <a href="#">property assignments</a>	An optional list of property assignments for the Relationship Template.
attributes	no	list of <a href="#">attribute assignments</a>	An optional list of attribute assignments for the Relationship Template.
interfaces	no	list of <a href="#">interface definitions</a>	An optional list of named interface definitions for the Node Template.
copy	no	<a href="#">string</a>	The optional (symbolic) name of another relationship template to copy into (all keynames and values) and use as a basis for this relationship template.

2026 **3.7.4.2 Grammar**

```

<relationship_template_name>:
  type: <relationship type name>
  description: <relationship type description>
  metadata:
    <map of string>
  properties:

```

```
<property_assignments>
attributes:
  <attribute_assignments>
interfaces:
  <interface_definitions>
copy:
  <source_relationship_template_name>
```

2027 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 2028 • **relationship\_template\_name**: represents the required symbolic name of the Relationship
- 2029 Template being declared.
- 2030 • **relationship\_type\_name**: represents the name of the Relationship Type the Relationship
- 2031 Template is based upon.
- 2032 • **relationship\_template\_description**: represents the optional [description](#) string for the
- 2033 Relationship Template.
- 2034 • **property\_assignments**: represents the optional list of [property assignments](#) for the Relationship
- 2035 Template that provide values for properties defined in its declared Relationship Type.
- 2036 • **attribute\_assignments**: represents the optional list of [attribute assignments](#) for the
- 2037 Relationship Template that provide values for attributes defined in its declared Relationship Type.
- 2038 • **interface\_definitions**: represents the optional list of [interface definitions](#) for the Relationship
- 2039 Template that augment those provided by its declared Relationship Type.
- 2040 • **source\_relationship\_template\_name**: represents the optional (symbolic) name of another
- 2041 relationship template to copy into (all keynames and values) and use as a basis for this
- 2042 relationship template.

### 2043 3.7.4.3 Additional requirements

- 2044 • The source relationship template provided as a value on the **copy** keyname MUST NOT itself use
- 2045 the **copy** keyname (i.e., it must itself be a complete relationship template description and not
- 2046 copied from another relationship template).

### 2047 3.7.4.4 Example

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

## 2048 3.7.5 Group definition

2049 A group definition defines a logical grouping of node templates, typically for management purposes, but is  
2050 separate from the application's topology template.

### 2051 3.7.5.1 Keynames

2052 The following is the list of recognized keynames for a TOSCA group definition:

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

### 2053 3.7.5.2 Grammar

2054 Group definitions have one the following grammars:

```

<group_name>:
  type: <group_type_name>
  description: <group_description>
  metadata:
    <map of string>
  properties:
    <property_assignments>
  members: [ <list_of_node_templates> ]
  interfaces:
    <interface_definitions>

```

2055 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 2056 • **group\_name**: represents the required symbolic name of the group as a [string](#).
- 2057 • **group\_type\_name**: represents the name of the Group Type the definition is based upon.
- 2058 • **group\_description**: contains an optional description of the group.
- 2059 • **property\_assignments**: represents the optional list of [property assignments](#) for the group
- 2060 definition that provide values for properties defined in its declared Group Type.
- 2061 • **list\_of\_node\_templates**: contains the required list of one or more node template names
- 2062 (within the same topology template) that are members of this logical group.
- 2063 • **interface\_definitions**: represents the optional list of [interface definitions](#) for the group
- 2064 definition that augment those provided by its declared Group Type.

### 2065 3.7.5.3 Additional Requirements

- 2066 • Group definitions **SHOULD NOT** be used to define or redefine relationships (dependencies) for
- 2067 an application that can be expressed using normative TOSCA Relationships within a TOSCA
- 2068 topology template.

2069 **3.7.5.4 Example**

2070 The following represents a group definition:

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

2071 **3.7.6 Policy definition**

2072 A policy definition defines a policy that can be associated with a TOSCA topology or top-level entity  
2073 definition (e.g., group definition, node template, etc.).

2074 **3.7.6.1 Keynames**

2075 The following is the list of recognized keynames for a TOSCA policy definition:

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

2076 **3.7.6.2 Grammar**

2077 Policy definitions have one the following grammars:

```
<policy\_name>:  
  type: <policy\_type\_name>  
  description: <policy\_description>  
  metadata:  
    <map of string>  
  properties:  
    <property assignments>  
  targets: [<list\_of\_policy\_targets>]  
  triggers:  
    <list\_of\_trigger\_definitions>
```

2078 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 2079 • **policy\_name**: represents the required symbolic name of the policy as a [string](#).
- 2080 • **policy\_type\_name**: represents the name of the policy the definition is based upon.

- 2081 • **policy\_description**: contains an optional description of the policy.
- 2082 • **property\_assignments**: represents the optional list of [property assignments](#) for the policy
- 2083 definition that provide values for properties defined in its declared Policy Type.
- 2084 • **list\_of\_policy\_targets**: represents the optional list of names of node templates or groups
- 2085 that the policy is to applied to.
- 2086 • **list\_of\_trigger\_definitions**: represents the optional list of [trigger definitions](#) for the policy.

### 2087 3.7.6.3 Example

2088 The following represents a policy definition:

```

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

```

## 2089 3.7.7 Imperative Workflow definition

2090 A workflow definition defines an imperative workflow that is associated with a TOSCA topology.

### 2091 3.7.7.1 Keynames

2092 The following is the list of recognized keynames for a TOSCA workflow definition:

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

### 2094 3.7.7.2 Grammar

2095 Imperative workflow definitions have the following grammar:

```

<workflow_name>:
  description: <workflow_description>
  metadata:
    <map of string>
  inputs:
    <property_definitions>

```

```

preconditions:
- <workflow_precondition_definition>
steps:
  <workflow_steps>

```

2096 In the above grammar, the pseudo values that appear in angle

## 2097 3.8 Topology Template definition

2098 This section defines the topology template of a cloud application. The main ingredients of the topology  
 2099 template are node templates representing components of the application and relationship templates  
 2100 representing links between the components. These elements are defined in the nested **node\_templates**  
 2101 section and the nested **relationship\_templates** sections, respectively. Furthermore, a topology  
 2102 template allows for defining input parameters, output parameters as well as grouping of node templates.

### 2103 3.8.1 Keynames

2104 The following is the list of recognized keynames for a TOSCA Topology Template:

Keyname	Required	Type	Description
description	no	<a href="#">description</a>	The optional description for the Topology Template.
inputs	no	list of <a href="#">parameter definitions</a>	An optional list of input parameters (i.e., as parameter definitions) for the Topology Template.
node_templates	no	list of <a href="#">node templates</a>	An optional list of node template definitions for the Topology Template.
relationship_templates	no	list of <a href="#">relationship templates</a>	An optional list of relationship templates for the Topology Template.
groups	no	list of <a href="#">group definitions</a>	An optional list of Group definitions whose members are node templates defined within this same Topology Template.
policies	no	list of <a href="#">policy definitions</a>	An optional list of Policy definitions for the Topology Template.
outputs	no	list of <a href="#">parameter definitions</a>	An optional list of output parameters (i.e., as parameter definitions) for the Topology Template.
substitution_mappings	no	N/A	An optional declaration that exports the topology template as an implementation of a Node type.  This also includes the mappings between the external Node Types named capabilities and requirements to existing implementations of those capabilities and requirements on Node templates declared within the topology template.

Keyname	Required	Type	Description
workflows	no	list of imperative workflow definitions	An optional map of imperative workflow definition for the Topology Template.

## 2105 3.8.2 Grammar

2106 The overall grammar of the **topology\_template** section is shown below.—Detailed grammar definitions  
2107 of the each sub-sections are provided in subsequent subsections.

```

topology_template:
  description: <template_description>
  inputs: <input_parameter_list>
  outputs: <output_parameter_list>
  node_templates: <node_template_list>
  relationship_templates: <relationship_template_list>
  groups: <group_definition_list>
  policies:
    - <policy_definition_list>
  workflows: <workflow_list>
  # Optional declaration that exports the Topology Template
  # as an implementation of a Node Type.
  substitution_mappings:
    node_type: <node_type_name>
  capabilities:
    <map_of_capability_mappings_to_expose>
  requirements:
    <map_of_requirement_mapping_to_expose>

```

2108 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 2109 • **template\_description**: represents the optional [description](#) string for Topology Template.
- 2110 • **input\_parameter\_list**: represents the optional list of input parameters (i.e., as property  
2111 definitions) for the Topology Template.
- 2112 • **output\_parameter\_list**: represents the optional list of output parameters (i.e., as property  
2113 definitions) for the Topology Template.
- 2114 • **group\_definition\_list**: represents the optional list of [group definitions](#) whose members are  
2115 node templates that also are defined within this Topology Template.
- 2116 • **policy\_definition\_list**: represents the optional sequenced list of policy definitions for the  
2117 Topology Template.
- 2118 • **workflow\_list**: represents the optional list of imperative workflow definitions  
2119 for the Topology Template.
- 2120 • **node\_template\_list**: represents the optional list of [node template](#) definitions for the Topology  
2121 Template.
- 2122 • **relationship\_template\_list**: represents the optional list of [relationship templates](#) for the  
2123 Topology Template.

- 2124 • **node\_type\_name**: represents the optional name of a [Node Type](#) that the Topology Template  
2125 implements as part of the **substitution\_mappings**.
- 2126 • **map\_of\_capability\_mappings\_to\_expose**: represents the mappings that expose internal  
2127 capabilities from node templates (within the topology template) as capabilities of the Node Type  
2128 definition that is declared as part of the **substitution\_mappings**.
- 2129 • **map\_of\_requirement\_mappings\_to\_expose**: represents the mappings of link requirements of  
2130 the Node Type definition that is declared as part of the **substitution\_mappings** to internal  
2131 requirements implementations within node templates (declared within the topology template).  
2132

2133 More detailed explanations for each of the Topology Template grammar's keynames appears in the  
2134 sections below.

### 2135 **3.8.2.1 inputs**

2136 The **inputs** section provides a means to define parameters using TOSCA parameter definitions, their  
2137 allowed values via constraints and default values within a TOSCA Simple Profile template. Input  
2138 parameters defined in the **inputs** section of a topology template can be mapped to properties of node  
2139 templates or relationship templates within the same topology template and can thus be used for  
2140 parameterizing the instantiation of the topology template.

2141  
2142 This section defines topology template-level input parameter section.

- 2143 • Inputs here would ideally be mapped to BoundaryDefinitions in TOSCA v1.0.
- 2144 • Treat input parameters as fixed global variables (not settable within template)
- 2145 • If not in input take default (nodes use default)

#### 2146 **3.8.2.1.1 Grammar**

2147 The grammar of the **inputs** section is as follows:

```
inputs:  
  <parameter_definition_list>
```

#### 2148 **3.8.2.1.2 Examples**

2149 This section provides a set of examples for the single elements of a topology template.

2150 Simple **inputs** example without any constraints:

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

2151 Example of **inputs** with constraints:

```
inputs:  
  SiteName:  
    type: string
```

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

## 2152 **3.8.2.2 node\_templates**

2153 The **node\_templates** section lists the Node Templates that describe the (software) components that are  
2154 used to compose cloud applications.

### 2155 **3.8.2.2.1 grammar**

2156 The grammar of the **node\_templates** section is as follows:

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

### 2157 **3.8.2.2.2 Example**

2158 Example of **node\_templates** section:

```
node_templates:
  my_webapp_node_template:
    type: WebApplication

  my_database_node_template:
    type: Database
```

## 2159 **3.8.2.3 relationship\_templates**

2160 The **relationship\_templates** section lists the Relationship Templates that describe the relations  
2161 between components that are used to compose cloud applications.

2162

2163 Note that in the TOSCA Simple Profile, the explicit definition of relationship templates as it was required  
2164 in TOSCA v1.0 is optional, since relationships between nodes get implicitly defined by referencing other  
2165 node templates in the requirements sections of node templates.

### 2166 **3.8.2.3.1 Grammar**

2167 The grammar of the **relationship\_templates** section is as follows:

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

2168 **3.8.2.3.2 Example**

2169 Example of `relationship_templates` section:

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

2170 **3.8.2.4 outputs**

2171 The **outputs** section provides a means to define the output parameters that are available from a TOSCA  
2172 Simple Profile service template. It allows for exposing attributes of node templates or relationship  
2173 templates within the containing **topology\_template** to users of a service.

2174 **3.8.2.4.1 Grammar**

2175 The grammar of the **outputs** section is as follows:

```
outputs:
  <parameter def list>
```

2176 **3.8.2.4.2 Example**

2177 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 ] }
```

2178 **3.8.2.5 groups**

2179 The **groups** section allows for grouping one or more node templates within a TOSCA Service Template  
2180 and for assigning special attributes like policies to the group.

2181 **3.8.2.5.1 Grammar**

2182 The grammar of the **groups** section is as follows:

```
groups:
  <group defn 1>
  ...
  <group defn n>
```

2183 **3.8.2.5.2 Example**

2184 The following example shows the definition of three Compute nodes in the **node\_templates** section of a  
2185 **topology\_template** as well as the grouping of two of the Compute nodes in a group **server\_group\_1**.

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

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

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

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

2186 **3.8.2.6 policies**

2187 The **policies** section allows for declaring policies that can be applied to entities in the topology template.

2188 **3.8.2.6.1 Grammar**

2189 The grammar of the **policies** section is as follows:

```
policies:
- <policy defn 1>
- ...
- <policy defn n>
```

2190 **3.8.2.6.2 Example**

2191 The following example shows the definition of a placement policy.

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

2192 **3.8.2.7 Notes**

- 2193 • The parameters (properties) that are listed as part of the **inputs** block can be mapped to
- 2194 **PropertyMappings** provided as part of **BoundaryDefinitions** as described by the TOSCA v1.0
- 2195 specification.
- 2196 • The node templates listed as part of the **node\_templates** block can be mapped to the list of
- 2197 **NodeTemplate** definitions provided as part of **TopologyTemplate** of a **ServiceTemplate** as
- 2198 described by the TOSCA v1.0 specification.
- 2199 • The relationship templates listed as part of the **relationship\_templates** block can be mapped
- 2200 to the list of **RelationshipTemplate** definitions provided as part of **TopologyTemplate** of a
- 2201 **ServiceTemplate** as described by the TOSCA v1.0 specification.
- 2202 • The output parameters that are listed as part of the **outputs** section of a topology template can
- 2203 be mapped to **PropertyMappings** provided as part of **BoundaryDefinitions** as described by
- 2204 the TOSCA v1.0 specification.
- 2205 ○ Note, however, that TOSCA v1.0 does not define a direction (input vs. output) for those
- 2206 mappings, i.e. TOSCA v1.0 **PropertyMappings** are underspecified in that respect and
- 2207 TOSCA Simple Profile's **inputs** and **outputs** provide a more concrete definition of input
- 2208 and output parameters.

2209 **3.9 Service Template definition**

2210 A TOSCA Service Template (YAML) document contains element definitions of building blocks for cloud  
 2211 application, or complete models of cloud applications. This section describes the top-level structural  
 2212 elements (TOSCA keynames) along with their grammars, which are allowed to appear in a TOSCA  
 2213 Service Template document.

2214 **3.9.1 Keynames**

2215 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.

Keyname	Required	Type	Description
artifact_types	no	list of <a href="#">Artifact Types</a>	This section contains an optional list of artifact type definitions for use in the service template
data_types	no	list of <a href="#">Data Types</a>	Declares a list of optional TOSCA Data Type definitions.
capability_types	no	list of <a href="#">Capability Types</a>	This section contains an optional list of capability type definitions for use in the service template.
interface_types	no	list of <a href="#">Interface Types</a>	This section contains an optional list of interface type definitions for use in the service template.
relationship_types	no	list of <a href="#">Relationship Types</a>	This section contains a set of relationship type definitions for use in the service template.
node_types	no	list of <a href="#">Node Types</a>	This section contains a set of node type definitions for use in the service template.
group_types	no	list of <a href="#">Group Types</a>	This section contains a list of group type definitions for use in the service template.
policy_types	no	list of <a href="#">Policy Types</a>	This section contains a list of policy type definitions for use in the service template.
topology_template	no	<a href="#">Topology Template definition</a>	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.

2216 **3.9.1.1 Metadata keynames**

2217 The following is the list of recognized metadata keynames for a TOSCA Service Template definition:

Keyname	Required	Type	Description
template_name	no	<a href="#">string</a>	Declares a descriptive name for the template.
template_author	no	<a href="#">string</a>	Declares the author(s) or owner of the template.
template_version	no	<a href="#">string</a>	Declares the version string for the template.

2218 **3.9.2 Grammar**

2219 The overall structure of a TOSCA Service Template and its top-level key collations using the TOSCA  
 2220 Simple Profile is shown below:

```
tosca_definitions_version: # Required TOSCA Definitions version string

# Optional metadata keyname: value pairs
metadata:
  template_name: <value>           # Optional name of this service template
  template_author: <value>         # Optional author of this service template
  template_version: <value>        # Optional version of this service template
# Optional list of domain or profile specific metadata keynames
```

```
# Optional description of the definitions inside the file.
description: <template type description>

dsl_definitions:
  # list of YAML alias anchors (or macros)

repositories:
  # list of external repository definitions which host TOSCA artifacts

imports:
  # ordered list of import definitions

artifact_types:
  # list of artifact type definitions

data_types:
  # list of datatype definitions

capability_types:
  # list of capability type definitions

interface_types
  # list of interface type definitions

relationship_types:
  # list of relationship type definitions

node_types:
  # list of node type definitions

group_types:
  # list of group type definitions

policy_types:
  # list of policy type definitions

topology_template:
  # topology template definition of the cloud application or service
```

2221 **3.9.2.1 Notes**

- 2222
- 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.
- 2223
- 2224

2225 **3.9.3 Top-level keyname definitions**

2226 **3.9.3.1 `tosca_definitions_version`**

2227 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.

2228

2229

2230 **3.9.3.1.1 Keyname**

```
tosca_definitions_version
```

2231 **3.9.3.1.2 Grammar**

2232 Single-line form:

```
tosca_definitions_version: <tosca_simple_profile_version>
```

2233 **3.9.3.1.3 Examples:**

2234 TOSCA Simple Profile version 1.0 specification using the defined namespace alias (see Section 3.1):

```
tosca_definitions_version: toska_simple_yaml_1_0
```

2235 TOSCA Simple Profile version 1.0 specification using the fully defined (target) namespace (see Section 3.1):

2236

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

2237 **3.9.3.2 metadata**

2238 This keyname is used to associate domain-specific metadata with the Service Template. The metadata keyname allows a declaration of a map of keynames with string values.

2239

2240 **3.9.3.2.1 Keyname**

```
metadata
```

2241 **3.9.3.2.2 Grammar**

```
metadata:  
  <map_of_string_values>
```

2242 **3.9.3.2.3 Example**

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

2243

2244 **3.9.3.3 template\_name**

2245 This optional metadata keyname can be used to declare the name of service template as a single-line  
2246 string value.

2247 **3.9.3.3.1 Keyname**

```
template_name
```

2248 **3.9.3.3.2 Grammar**

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

2249 **3.9.3.3.3 Example**

```
template_name: My service template
```

2250 **3.9.3.3.4 Notes**

- 2251
- Some service templates are designed to be referenced and reused by other service templates.  
2252 Therefore, in these cases, the **template\_name** value SHOULD be designed to be used as a  
2253 unique identifier through the use of namespacing techniques.

2254 **3.9.3.4 template\_author**

2255 This optional metadata keyname can be used to declare the author(s) of the service template as a single-  
2256 line string value.

2257 **3.9.3.4.1 Keyname**

```
template_author
```

2258 **3.9.3.4.2 Grammar**

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

2259 **3.9.3.4.3 Example**

```
template_author: My service template
```

2260 **3.9.3.5 template\_version**

2261 This optional metadata keyname can be used to declare a domain specific version of the service template  
2262 as a single-line string value.

2263 **3.9.3.5.1 Keyname**

```
template_version
```

2264 **3.9.3.5.2 Grammar**

```
template_version: <version>
```

2265 **3.9.3.5.3 Example**

```
template_version: 2.0.17
```

2266 **3.9.3.5.4 Notes:**

- 2267
- Some service templates are designed to be referenced and reused by other service templates  
2268 and have a lifecycle of their own. Therefore, in these cases, a **template\_version** value  
2269 SHOULD be included and used in conjunction with a unique **template\_name** value to enable  
2270 lifecycle management of the service template and its contents.

2271 **3.9.3.6 description**

2272 This optional keyname provides a means to include single or multiline descriptions within a TOSCA  
2273 Simple Profile template as a scalar string value.

2274 **3.9.3.6.1 Keyname**

```
description
```

2275 **3.9.3.7 dsl\_definitions**

2276 This optional keyname provides a section to define macros (e.g., YAML-style macros when using the  
2277 TOSCA Simple Profile in YAML specification).

2278 **3.9.3.7.1 Keyname**

```
dsl_definitions
```

2279 **3.9.3.7.2 Grammar**

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

2280 **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
```

2281 **3.9.3.8 repositories**

2282 This optional keyname provides a section to define external repositories which may contain artifacts or  
2283 other TOSCA Service Templates which might be referenced or imported by the TOSCA Service Template  
2284 definition.

2285 **3.9.3.8.1 Keyname**

```
repositories
```

2286 **3.9.3.8.2 Grammar**

```
repositories:
  <repository definition 1>
  ...
  <repository definition n>
```

2287 **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/
```

2288 **3.9.3.9 imports**

2289 This optional keyname provides a way to import a *block sequence* of one or more TOSCA Definitions  
2290 documents. TOSCA Definitions documents can contain reusable TOSCA type definitions (e.g., Node  
2291 Types, Relationship Types, Artifact Types, etc.) defined by other authors. This mechanism provides an  
2292 effective way for companies and organizations to define normative types and/or describe their software  
2293 applications for reuse in other TOSCA Service Templates.

2294 **3.9.3.9.1 Keyname**

```
imports
```

2295 **3.9.3.9.2 Grammar**

```
imports:  
  - <import definition 1>  
  - ...  
  - <import definition n>
```

2296 **3.9.3.9.3 Example**

```
# An example import of definitions files from a location relative to the  
# file location of the service template declaring the import.  
imports:  
  - some_definitions: relative_path/my_defns/my_typesdefs_1.yaml  
  - file: my_defns/my_typesdefs_n.yaml  
    repository: my_company_repo  
    namespace_uri: http://mycompany.com/ns/tosca/2.0  
    namespace_prefix: mycompany
```

2297 **3.9.3.10 artifact\_types**

2298 This optional keyname lists the Artifact Types that are defined by this Service Template.

2299 **3.9.3.10.1 Keyname**

```
artifact_types
```

2300 **3.9.3.10.2 Grammar**

```
artifact_types:  
  <artifact type defn 1>  
  ...  
  <artifact type defn n>
```

2301 **3.9.3.10.3 Example**

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

2302 **3.9.3.11 data\_types**

2303 This optional keyname provides a section to define new data types in TOSCA.

2304 **3.9.3.11.1 Keyname**

```
data_types
```

2305 **3.9.3.11.2 Grammar**

```
data_types:  
  <tosca_datatype_def 1>  
  ...  
  <tosca_datatype_def n>
```

2306 **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
```

2307 **3.9.3.12 capability\_types**

2308 This optional keyname lists the Capability Types that provide the reusable type definitions that can be  
2309 used to describe features Node Templates or Node Types can declare they support.

### 2310 3.9.3.12.1 Keyname

```
capability_types
```

### 2311 3.9.3.12.2 Grammar

```
capability_types:  
  <capability_type_defn 1>  
  ...  
  <capability_type_defn n>
```

### 2312 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 ...
```

### 2313 3.9.3.13 interface\_types

2314 This optional keyname lists the Interface Types that provide the reusable type definitions that can be used  
2315 to describe operations for on TOSCA entities such as Relationship Types and Node Types.

### 2316 3.9.3.13.1 Keyname

```
interface_types
```

### 2317 3.9.3.13.2 Grammar

```
interface_types:  
  <interface_type_defn 1>  
  ...  
  <interface_type_defn n>
```

### 2318 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.
```

### 2319 **3.9.3.14 relationship\_types**

2320 This optional keyname lists the Relationship Types that provide the reusable type definitions that can be  
2321 used to describe dependent relationships between Node Templates or Node Types.

#### 2322 **3.9.3.14.1 Keyname**

```
relationship_types
```

#### 2323 **3.9.3.14.2 Grammar**

```
relationship_types:
  <relationship type defn 1>
  ...
  <relationship type defn n>
```

#### 2324 **3.9.3.14.3 Example**

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

  mycompany.mytypes.myCustomConnectionType:
    derived_from: toska.relationships.ConnectsTo
    properties:
      # more details ...
```

### 2325 **3.9.3.15 node\_types**

2326 This optional keyname lists the Node Types that provide the reusable type definitions for software  
2327 components that Node Templates can be based upon.

#### 2328 **3.9.3.15.1 Keyname**

```
node_types
```

#### 2329 **3.9.3.15.2 Grammar**

```
node_types:
  <node type defn 1>
```

```
...  
<node_type_defn n>
```

### 2330 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
```

### 2331 3.9.3.15.4 Notes

- 2332
- 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.
- 2333

### 2334 3.9.3.16 group\_types

2335 This optional keyname lists the Group Types that are defined by this Service Template.

#### 2336 3.9.3.16.1 Keyname

```
group_types
```

#### 2337 3.9.3.16.2 Grammar

```
group_types:  
  <group_type_defn 1>  
  ...  
  <group_type_defn n>
```

### 2338 3.9.3.16.3 Example

```
group_types:  
  mycompany.mtypes.myScalingGroup:  
    derived_from: tosca.groups.Root
```

### 2339 3.9.3.17 policy\_types

2340 This optional keyname lists the Policy Types that are defined by this Service Template.

2341 **3.9.3.17.1 Keyname**

```
policy_types
```

2342 **3.9.3.17.2 Grammar**

```
policy_types:  
  <policy type defn 1>  
  ...  
  <policy type defn n>
```

2343 **3.9.3.17.3 Example**

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

2344

## 4 TOSCA functions

2345 Except for the examples, this section is **normative** and includes functions that are supported for use  
2346 within a TOSCA Service Template.

### 4.1 Reserved Function Keywords

2348 The following keywords MAY be used in some TOSCA function in place of a TOSCA Node or  
2349 Relationship Template name. A TOSCA orchestrator will interpret them at the time the function would be  
2350 evaluated at runtime as described in the table below. Note that some keywords are only valid in the  
2351 context of a certain TOSCA entity as also denoted in the table.

2352

Keyword	Valid Contexts	Description
SELF	Node Template or Relationship Template	A TOSCA orchestrator will interpret this keyword as the Node or Relationship Template instance that contains the function at the time the function is evaluated.
SOURCE	Relationship Template only.	A TOSCA orchestrator will interpret this keyword as the Node Template instance that is at the source end of the relationship that contains the referencing function.
TARGET	Relationship Template only.	A TOSCA orchestrator will interpret this keyword as the Node Template instance that is at the target end of the relationship that contains the referencing function.
HOST	Node Template only	A TOSCA orchestrator will interpret this keyword to refer to the all nodes that "host" the node using this reference (i.e., as identified by its HostedOn relationship).  Specifically, TOSCA orchestrators that encounter this keyword when evaluating the <b>get_attribute</b> or <b>get_property</b> 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.

2353

### 4.2 Environment Variable Conventions

#### 4.2.1 Reserved Environment Variable Names and Usage

2356 TOSCA orchestrators utilize certain reserved keywords in the execution environments that  
2357 implementation artifacts for Node or Relationship Templates operations are executed in. They are used to  
2358 provide information to these implementation artifacts such as the results of TOSCA function evaluation or  
2359 information about the instance model of the TOSCA application

2360

2361 The following keywords are reserved environment variable names in any TOSCA supported execution  
2362 environment:

Keyword	Valid Contexts	Description
TARGETS	Relationship Template only.	<ul style="list-style-type: none"> <li>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.</li> <li>The value of this environment variable will be a comma-separated list of identifiers of the single target node instances (i.e., the <b>tosca_id</b> attribute of the node).</li> </ul>
TARGET	Relationship Template only.	<ul style="list-style-type: none"> <li>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.</li> <li>The value of this environment variable will be the identifier of the single target node instance (i.e., the <b>tosca_id</b> attribute of the node).</li> </ul>
SOURCES	Relationship Template only.	<ul style="list-style-type: none"> <li>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.</li> <li>The value of this environment variable will be a comma-separated list of identifiers of the single source node instances (i.e., the <b>tosca_id</b> attribute of the node).</li> </ul>
SOURCE	Relationship Template only.	<ul style="list-style-type: none"> <li>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.</li> <li>The value of this environment variable will be the identifier of the single source node instance (i.e., the <b>tosca_id</b> attribute of the node).</li> </ul>

2363

2364 For scripts (or implementation artifacts in general) that run in the context of relationship operations, select  
 2365 properties and attributes of both the relationship itself as well as select properties and attributes of the  
 2366 source and target node(s) of the relationship can be provided to the environment by declaring respective  
 2367 operation inputs.

2368

2369 Declared inputs from mapped properties or attributes of the source or target node (selected via the  
 2370 **SOURCE** or **TARGET** keyword) will be provided to the environment as variables having the exact same name  
 2371 as the inputs. In addition, the same values will be provided for the complete set of source or target nodes,  
 2372 however prefixed with the ID if the respective nodes. By means of the **SOURCES** or **TARGETS** variables  
 2373 holding the complete set of source or target node IDs, scripts will be able to iterate over corresponding  
 2374 inputs for each provided ID prefix.

2375

2376 The following example snippet shows an imaginary relationship definition from a load-balancer node to  
 2377 worker nodes. A script is defined for the **add\_target** operation of the Configure interface of the  
 2378 relationship, and the **ip\_address** attribute of the target is specified as input to the script:

2379

```
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

```

2380 The **add\_target** operation will be invoked, whenever a new target member is being added to the load-  
 2381 balancer. With the above inputs declaration, a **member\_ip** environment variable that will hold the IP  
 2382 address of the target being added will be provided to the **configure\_members.py** script. In addition, the  
 2383 IP addresses of all current load-balancer members will be provided as environment variables with a  
 2384 naming scheme of **<target node ID>\_member\_ip**. This will allow, for example, scripts that always just  
 2385 write the complete list of load-balancer members into a configuration file to do so instead of updating  
 2386 existing list, which might be more complicated.

2387 Assuming that the TOSCA application instance includes five load-balancer members, **node1** through  
 2388 **node5**, where **node5** is the current target being added, the following environment variables (plus  
 2389 potentially more variables) would be provided to the script:

```

# the ID of the current target and the IDs of all targets
TARGET=node5
TARGETS=node1,node2,node3,node4,node5

# the input for the current target and the inputs of all targets
member_ip=10.0.0.5
node1_member_ip=10.0.0.1
node2_member_ip=10.0.0.2
node3_member_ip=10.0.0.3
node4_member_ip=10.0.0.4
node5_member_ip=10.0.0.5

```

2390 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 ...

```

2391 **4.2.2 Prefixed vs. Unprefixed TARGET names**

2392 The list target node types assigned to the TARGETS key in an execution environment would have names  
2393 prefixed by unique IDs that distinguish different instances of a node in a running model Future drafts of  
2394 this specification will show examples of how these names/IDs will be expressed.

2395 **4.2.2.1 Notes**

- 2396 • Target of interest is always un-prefixed. Prefix is the target opaque ID. The IDs can be used to  
2397 find the environment var. for the corresponding target. Need an example here.
- 2398 • If you have one node that contains multiple targets this would also be used (add or remove target  
2399 operations would also use this you would get set of all current targets).

2400 **4.3 Intrinsic functions**

2401 These functions are supported within the TOSCA template for manipulation of template data.

2402 **4.3.1 concat**

2403 The **concat** function is used to concatenate two or more string values within a TOSCA service template.

2404 **4.3.1.1 Grammar**

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

2405 **4.3.1.2 Parameters**

Parameter	Required	Type	Description
<string_value_expressions_*>	yes	list of <a href="#">string</a> or <a href="#">string</a> 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.

2406 **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 ] ] }
```

2407 **4.3.2 token**

2408 The **token** function is used within a TOSCA service template on a string to parse out (tokenize)  
2409 substrings separated by one or more token characters within a larger string.

2410 **4.3.2.1 Grammar**

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

2411 **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.

2412 **4.3.2.3 Examples**

```

outputs:
  webservice_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 ] }

```

2413 **4.4 Property functions**

2414 These functions are used within a service template to obtain property values from property definitions  
 2415 declared elsewhere in the same service template. These property definitions can appear either directly in  
 2416 the service template itself (e.g., in the inputs section) or on entities (e.g., node or relationship templates)  
 2417 that have been modeled within the template.

2418  
 2419 Note that the **get\_input** and **get\_property** functions may only retrieve the static values of property  
 2420 definitions of a TOSCA application as defined in the TOSCA Service Template. The **get\_attribute**  
 2421 function should be used to retrieve values for attribute definitions (or property definitions reflected as  
 2422 attribute definitions) from the runtime instance model of the TOSCA application (as realized by the  
 2423 TOSCA orchestrator).

2424 **4.4.1 get\_input**

2425 The **get\_input** function is used to retrieve the values of properties declared within the **inputs** section of  
 2426 a TOSCA Service Template.

2427 **4.4.1.1 Grammar**

```

get_input: <input_property_name>

```

2428 **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.

2429 **4.4.1.3 Examples**

```

inputs:
  cpus:
    type: integer

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

```

2430 **4.4.2 get\_property**

2431 The `get_property` function is used to retrieve property values between modelable entities defined in the  
 2432 same service template.

2433 **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> ]

```

2434 **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.  <b>Note:</b> 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 <b>list</b> 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.

2435 **4.4.2.3 Examples**

2436 The following example shows how to use the `get_property` function with an actual Node Template  
2437 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 ] }
```

2438 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 ] }

```

2439 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 ] }

```

## 2440 4.5 Attribute functions

2441 These functions (attribute functions) are used within an instance model to obtain attribute values from  
 2442 instances of nodes and relationships that have been created from an application model described in a  
 2443 service template. The instances of nodes or relationships can be referenced by their name as assigned  
 2444 in the service template or relative to the context where they are being invoked.

### 2445 4.5.1 get\_attribute

2446 The `get_attribute` function is used to retrieve the values of named attributes declared by the  
 2447 referenced node or relationship template name.

#### 2448 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> ]

```

#### 2449 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.  <b>Note:</b> 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.

Parameter	Required	Type	Description
<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 <b>list</b> 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.

2450 **4.5.1.3 Examples:**

2451 The attribute functions are used in the same way as the equivalent Property functions described above.  
2452 Please see their examples and replace “get\_property” with “get\_attribute” function name.

2453 **4.5.1.4 Notes**

2454 These functions are used to obtain attributes from instances of node or relationship templates by the  
2455 names they were given within the service template that described the application model (pattern).

- 2456 • These functions only work when the orchestrator can resolve to a single node or relationship  
2457 instance for the named node or relationship. This essentially means this is acknowledged to work  
2458 only when the node or relationship template being referenced from the service template has a  
2459 cardinality of 1 (i.e., there can only be one instance of it running).

2460 **4.6 Operation functions**

2461 These functions are used within an instance model to obtain values from interface operations. These can  
2462 be used in order to set an attribute of a node instance at runtime or to pass values from one operation to  
2463 another.

2464 **4.6.1 get\_operation\_output**

2465 The `get_operation_output` function is used to retrieve the values of variables exposed / exported from  
2466 an interface operation.

2467 **4.6.1.1 Grammar**

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

2468 **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.

2469 **4.6.1.3 Notes**

- If operation failed, then ignore its outputs. Orchestrators should allow orchestrators to continue running when possible past deployment in the lifecycle. For example, if an update fails, the application should be allowed to continue running and some other method would be used to alert administrators of the failure.

2474 **4.7 Navigation functions**

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

2476 **4.7.1 get\_nodes\_of\_type**

2477 The `get_nodes_of_type` function can be used to retrieve a list of all known instances of nodes of the  
2478 declared Node Type.

2479 **4.7.1.1 Grammar**

```
get_nodes_of_type: <node_type_name>
```

2480 **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.

2481 **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 <code>node_type_name</code> supplied as an input parameter of this function.

2482 **4.8 Artifact functions**

2483 **4.8.1 get\_artifact**

2484 The `get_artifact` function is used to retrieve artifact location between modelable entities defined in the  
2485 same service template.

2486 **4.8.1.1 Grammar**

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

2487 **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.

Parameter	Required	Type	Description
<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 <b>get_artifact</b> 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 <b>get_artifact</b> 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>

2488 **4.8.1.3 Examples**

2489 The following example uses a snippet of a WordPress [WordPress] web application to show how to use  
 2490 the **get\_artifact** function with an actual Node Template name:

2491 **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
```

2492 In such implementation the TOSCA orchestrator may provide the **wordpress.zip** archive as

- 2493
- a local URL (example: <file:///home/user/wordpress.zip>) or
  - a remote one (example: <http://cloudrepo:80/files/wordpress.zip>) where some orchestrator may indeed provide some global artifact repository management features.
- 2494
- 2495

#### 2496 **4.8.1.3.2 Example: Retrieving artifact as a local path**

2497 The following example explains how to force the orchestrator to copy the file locally before calling the  
2498 operation's implementation script:

2499

```
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
```

2500 In such implementation the TOSCA orchestrator must provide the wordpress.zip archive as a local path  
2501 (example: <tmp/wordpress.zip>) and **will remove it** after the operation is completed.

#### 2502 **4.8.1.3.3 Example: Retrieving artifact in a specified location**

2503 The following example explains how to force the orchestrator to copy the file locally to a specific location  
2504 before calling the operation's implementation script :

2505

```
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
```

2506 In such implementation the TOSCA orchestrator must provide the wordpress.zip archive as a local path  
2507 (example: C:/wpdata/wp.zip ) and **will let it** after the operation is completed.

## 2508 **4.9 Context-based Entity names (global)**

2509 Future versions of this specification will address methods to access entity names based upon the context  
2510 in which they are declared or defined.

### 2511 **4.9.1.1 Goals**

- 2512 • Using the full paths of modelable entity names to qualify context with the future goal of a more  
2513 robust `get_attribute` function: e.g., `get_attribute( <context-based-entity-name>, <attribute name>)`

2514

## 5 TOSCA normative type definitions

2515  
2516

Except for the examples, this section is **normative** and contains normative type definitions which must be supported for conformance to this specification.

2517

2518  
2519  
2520  
2521

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

2522

### 5.1 Assumptions

2523  
2524  
2525  
2526  
2527

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

2528

### 5.2 TOSCA normative type names

2529

Every normative type has three names declared:

2530  
2531  
2532  
2533  
2534  
2535  
2536  
2537  
2538  
2539  
2540  
2541  
2542  
2543

1. **Type URI** – This is the unique identifying name for the type.
  - a. These are reserved names within the TOSCA namespace.
2. **Shorthand Name** – This is the shorter (simpler) name that can be used in place of its corresponding, full **Type URI** name.
  - a. These are reserved names within TOSCA namespace that MAY be used in place of the full Type URI.
  - b. Profiles of the OASIS TOSCA Simple Profile specification SHALL assure non-collision of names for new types when they are introduced.
  - c. TOSCA type designers SHOULD NOT create new types with names that would collide with any TOSCA normative type Shorthand Name.
3. **Type Qualified Name** – This is a modified **Shorthand Name** that includes the “**tosca:**” namespace prefix which clearly qualifies it as being part of the TOSCA namespace.
  - a. This name MAY be used to assure there is no collision when types are imported from other (non) TOSCA approved sources.

2544

#### 5.2.1 Additional requirements

2545  
2546  
2547  
2548  
2549

- **Case sensitivity** - TOSCA Type URI, Shorthand and Type Qualified names SHALL be treated as case sensitive.
  - The case of each type name has been carefully selected by the TOSCA working group and TOSCA orchestrators and processors SHALL strictly recognize the name casing as specified in this specification or any of its approved profiles.

2550

### 5.3 Data Types

2551

#### 5.3.1 `tosca.datatypes.Root`

2552

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

2553 **5.3.1.1 Definition**

2554 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
```

2555 **5.3.2 tosca.datatypes.Credential**

2556 The Credential type is a complex TOSCA data Type used when describing authorization credentials used  
2557 to access network accessible resources.

<b>Shorthand Name</b>	Credential
<b>Type Qualified Name</b>	tosca:Credential
<b>Type URI</b>	tosca.datatypes.Credential

2558 **5.3.2.1 Properties**

Name	Required	Type	Constraints	Description
protocol	no	string	None	The optional protocol name.
token_type	yes	string	default: password	The required token type.
token	yes	string	None	The required token used as a credential for authorization or access to a networked resource.
keys	no	map of string	None	The optional list of protocol-specific keys or assertions.
user	no	string	None	The optional user (name or ID) used for non-token based credentials.

2559 **5.3.2.2 Definition**

2560 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
```

### 2561 5.3.2.3 Additional requirements

- 2562
- TOSCA Orchestrators SHALL interpret and validate the value of the **token** property based upon the value of the **token\_type** property.
- 2563

### 2564 5.3.2.4 Notes

- 2565
- Specific token types and encoding them using network protocols are not defined or covered in this specification.
- 2566
- The use of transparent user names (IDs) or passwords are not considered best practice.
- 2567

### 2568 5.3.2.5 Examples

#### 2569 5.3.2.5.1 Provide a simple user name and password without a protocol or 2570 standardized token format

```
<some_tosca_entity>:
  properties:
    my_credential:
      type: Credential
      properties:
        user: myusername
        token: mypassword
```

#### 2571 5.3.2.5.2 HTTP Basic access authentication credential

```
<some_tosca_entity>:
  properties:
    my_credential: # type: Credential
      protocol: http
      token_type: basic_auth
      # Username and password are combined into a string
      # Note: this would be base64 encoded before transmission by any impl.
      token: myusername:mypassword
```

#### 2572 5.3.2.5.3 X-Auth-Token credential

```
<some_tosca_entity>:
```

```

properties:
  my_credential: # type: Credential
    protocol: xauth
    token_type: X-Auth-Token
    # token encoded in Base64
    token: 604bbe45ac7143a79e14f3158df67091

```

2573 **5.3.2.5.4 OAuth bearer token credential**

```

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

```

2574 **5.3.2.6 OpenStack SSH Keypair**

```

<some_tosca_entity>:
  properties:
    my_ssh_keypair: # type: Credential
      protocol: ssh
      token_type: identifier
      # token is a reference (ID) to an existing keypair (already installed)
      token: <keypair_id>

```

2575

2576 **5.3.3 toska.datatypes.TimeInterval**

2577 The TimeInterval type is a complex TOSCA data Type used when describing a period of time using the  
 2578 YAML ISO 8601 format to declare the start and end times.

<b>Shorthand Name</b>	TimeInterval
<b>Type Qualified Name</b>	tosca:TimeInterval
<b>Type URI</b>	tosca.datatypes.TimeInterval

2579 **5.3.3.1 Properties**

Name	Required	Type	Constraints	Description
start_time	yes	timestamp	None	The <b>inclusive</b> start time for the time interval.
end_time	yes	timestamp	None	The <b>inclusive</b> end time for the time interval.

2580 **5.3.3.2 Definition**

2581 The TOSCA TimeInterval type is defined as follows:

```
tosca.datatypes.TimeInterval:
  derived_from: toska.datatypes.Root
  properties:
    start_time:
      type: timestamp
      required: true
    end_time:
      type: timestamp
      required: true
```

2582 **5.3.3.3 Examples**

2583 **5.3.3.3.1 Multi-day evaluation time period**

```
properties:
  description:
    evaluation_period: Evaluate a service for a 5-day period across time zones
    type: TimeInterval
    start_time: 2016-04-04-15T00:00:00Z
    end_time: 2016-04-08T21:59:43.10-06:00
```

2584 **5.3.4 toska.datatypes.network.NetworkInfo**

2585 The Network type is a complex TOSCA data type used to describe logical network information.

<b>Shorthand Name</b>	NetworkInfo
<b>Type Qualified Name</b>	tosca:NetworkInfo
<b>Type URI</b>	tosca.datatypes.network.NetworkInfo

2586 **5.3.4.1 Properties**

Name	Type	Constraints	Description
network_name	string	None	The name of the logical network. e.g., "public", "private", "admin". etc.
network_id	string	None	The unique ID of for the network generated by the network provider.
addresses	string []	None	The list of IP addresses assigned from the underlying network.

2587 **5.3.4.2 Definition**

2588 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

```

### 2589 5.3.4.3 Examples

2590 Example usage of the NetworkInfo data type:

```

<some_tosca_entity>:
  properties:
    private_network:
      network_name: private
      network_id: 3e54214f-5c09-1bc9-9999-44100326da1b
      addresses: [ 10.111.128.10 ]

```

### 2591 5.3.4.4 Additional Requirements

- 2592 • It is expected that TOSCA orchestrators MUST be able to map the **network\_name** from the
- 2593 TOSCA model to underlying network model of the provider.
- 2594 • The properties (or attributes) of NetworkInfo may or may not be required depending on usage
- 2595 context.

### 2596 5.3.5 toska.datatypes.network.PortInfo

2597 The PortInfo type is a complex TOSCA data type used to describe network port information.

<b>Shorthand Name</b>	PortInfo
<b>Type Qualified Name</b>	tosca:PortInfo
<b>Type URI</b>	tosca.datatypes.network.PortInfo

### 2598 5.3.5.1 Properties

Name	Type	Constraints	Description
port_name	<a href="#">string</a>	None	The logical network port name.
port_id	<a href="#">string</a>	None	The unique ID for the network port generated by the network provider.
network_id	<a href="#">string</a>	None	The unique ID for the network.

Name	Type	Constraints	Description
mac_address	<a href="#">string</a>	None	The unique media access control address ( <b>MAC address</b> ) assigned to the port.
addresses	<a href="#">string</a> []	None	The list of IP address(es) assigned to the port.

2599 **5.3.5.2 Definition**

2600 The TOSCA PortInfo type is defined as follows:

```
tosca.datatypes.network.PortInfo:
  derived_from: tosca.datatypes.Root
  properties:
    port_name:
      type: string
    port_id:
      type: string
    network_id:
      type: string
    mac_address:
      type: string
    addresses:
      type: list
    entry_schema:
      type: string
```

2601 **5.3.5.3 Examples**

2602 Example usage of the PortInfo data type:

```
ethernet_port:
  properties:
    port_name: port1
    port_id: 2c0c7a37-691a-23a6-7709-2d10ad041467
    network_id: 3e54214f-5c09-1bc9-9999-44100326da1b
    mac_address: f1:18:3b:41:92:1e
    addresses: [ 172.24.9.102 ]
```

2603 **5.3.5.4 Additional Requirements**

- 2604 • It is expected that TOSCA orchestrators **MUST** be able to map the **port\_name** from the TOSCA
- 2605 • The properties (or attributes) of PortInfo may or may not be required depending on usage context.

2607 **5.3.6 tosca.datatypes.network.PortDef**

2608 The PortDef type is a TOSCA data Type used to define a network port.

<b>Shorthand Name</b>	PortDef
<b>Type Qualified Name</b>	tosca:PortDef
<b>Type URI</b>	tosca.datatypes.network.PortDef

2609 **5.3.6.1 Definition**

2610 The TOSCA PortDef type is defined as follows:

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

2611 **5.3.6.2 Examples**

2612 Simple usage of a PortDef property type:

```
properties:
  listen_port: 9090
```

2613 Example declaration of a property for a custom type based upon PortDef:

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

2614 **5.3.7 tosca.datatypes.network.PortSpec**

2615 The PortSpec type is a complex TOSCA data Type used when describing port specifications for a  
2616 network connection.

<b>Shorthand Name</b>	PortSpec
<b>Type Qualified Name</b>	tosca:PortSpec
<b>Type URI</b>	tosca.datatypes.network.PortSpec

2617 **5.3.7.1 Properties**

Name	Required	Type	Constraints	Description
protocol	yes	string	default: tcp	The required protocol used on the port.
source	no	PortDef	See PortDef	The optional source port.
source_range	no	range	in_range: [ 1, 65536 ]	The optional range for source port.
target	no	PortDef	See PortDef	The optional target port.
target_range	no	range	in_range: [ 1, 65536 ]	The optional range for target port.

2618 **5.3.7.2 Definition**

2619 The TOSCA PortSpec type is defined as follows:

```
tosca.datatypes.network.PortSpec:
  derived_from: tosca.datatypes.Root
  properties:
    protocol:
      type: string
      required: true
      default: tcp
      constraints:
        - valid_values: [ udp, tcp, igmp ]
    target:
      type: PortDef
      required: false
    target_range:
      type: range
      required: false
      constraints:
        - in_range: [ 1, 65535 ]
    source:
      type: PortDef
      required: false
    source_range:
      type: range
      required: false
      constraints:
        - in_range: [ 1, 65535 ]
```

2620 **5.3.7.3 Additional requirements**

- 2621
- 2622
- 2623
- 2624
- 2625
- 2626
- A valid PortSpec MUST have at least one of the following properties: **target**, **target\_range**, **source** or **source\_range**.
  - A valid PortSpec MUST have a value for the **source** property that is within the numeric range specified by the property **source\_range** when **source\_range** is specified.
  - A valid PortSpec MUST have a value for the **target** property that is within the numeric range specified by the property **target\_range** when **target\_range** is specified.

2627 **5.3.7.4 Examples**

2628 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 ]
```

## 2629 5.4 Artifact Types

2630 TOSCA Artifacts Types represent the types of packages and files used by the orchestrator when  
2631 deploying TOSCA Node or Relationship Types or invoking their interfaces. Currently, artifacts are  
2632 logically divided into three categories:

2633

- 2634 • **Deployment Types:** includes those artifacts that are used during deployment (e.g., referenced  
2635 on create and install operations) and include packaging files such as RPMs, ZIPs, or TAR files.
- 2636 • **Implementation Types:** includes those artifacts that represent imperative logic and are used to  
2637 implement TOSCA Interface operations. These typically include scripting languages such as  
2638 Bash (.sh), Chef [[Chef](#)] and Puppet [[Puppet](#)].
- 2639 • **Runtime Types:** includes those artifacts that are used during runtime by a service or component  
2640 of the application. This could include a library or language runtime that is needed by an  
2641 application such as a PHP or Java library.

2642

2643 **Note:** Additional TOSCA Artifact Types will be developed in future drafts of this specification.

### 2644 5.4.1 `tosca.artifacts.Root`

2645 This is the default (root) TOSCA [Artifact Type](#) definition that all other TOSCA base Artifact Types derive  
2646 from.

#### 2647 5.4.1.1 Definition

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

### 2648 5.4.2 `tosca.artifacts.File`

2649 This artifact type is used when an artifact definition needs to have its associated file simply treated as a  
2650 file and no special handling/handlers are invoked (i.e., it is not treated as either an implementation or  
2651 deployment artifact type).

<b>Shorthand Name</b>	File
<b>Type Qualified Name</b>	tosca:File
<b>Type URI</b>	tosca.artifacts.File

2652 **5.4.2.1 Definition**

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

2653 **5.4.3 Deployment Types**

2654 **5.4.3.1 tosca.artifacts.Deployment**

2655 This artifact type represents the parent type for all deployment artifacts in TOSCA. This class of artifacts  
 2656 typically represents a binary packaging of an application or service that is used to install/create or deploy  
 2657 it as part of a node's lifecycle.

2658 **5.4.3.1.1 Definition**

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

2659 **5.4.3.2 Additional Requirements**

- 2660 • TOSCA Orchestrators MAY throw an error if it encounters a non-normative deployment artifact  
 2661 type that it is not able to process.

2662 **5.4.3.3 tosca.artifacts.Deployment.Image**

2663 This artifact type represents a parent type for any "image" which is an opaque packaging of a TOSCA  
 2664 Node's deployment (whether real or virtual) whose contents are typically already installed and pre-  
 2665 configured (i.e., "stateful") and prepared to be run on a known target container.

<b>Shorthand Name</b>	Deployment.Image
<b>Type Qualified Name</b>	tosca:Deployment.Image
<b>Type URI</b>	tosca.artifacts.Deployment.Image

2666 **5.4.3.3.1 Definition**

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

2667 **5.4.3.4 tosca.artifacts.Deployment.Image.VM**

2668 This artifact represents the parent type for all Virtual Machine (VM) image and container formatted  
 2669 deployment artifacts. These images contain a stateful capture of a machine (e.g., server) including

2670 operating system and installed software along with any configurations and can be run on another  
2671 machine using a hypervisor which virtualizes typical server (i.e., hardware) resources.

#### 2672 5.4.3.4.1 Definition

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

#### 2673 5.4.3.4.2 Notes

- 2674
- Future drafts of this specification may include popular standard VM disk image (e.g., ISO, VMI, VMDX, QCOW2, etc.) and container (e.g., OVF, bare, etc.) formats. These would include consideration of disk formats such as:
- 2675  
2676

### 2677 5.4.4 Implementation Types

#### 2678 5.4.4.1 `tosca.artifacts.Implementation`

2679 This artifact type represents the parent type for all implementation artifacts in TOSCA. These artifacts are  
2680 used to implement operations of TOSCA interfaces either directly (e.g., scripts) or indirectly (e.g., config.  
2681 files).

#### 2682 5.4.4.1.1 Definition

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

#### 2683 5.4.4.2 Additional Requirements

- 2684
- TOSCA Orchestrators **MAY** throw an error if it encounters a non-normative implementation  
2685 artifact type that it is not able to process.

#### 2686 5.4.4.3 `tosca.artifacts.Implementation.Bash`

2687 This artifact type represents a Bash script type that contains Bash commands that can be executed on  
2688 the Unix Bash shell.

<b>Shorthand Name</b>	Bash
<b>Type Qualified Name</b>	tosca:Bash
<b>Type URI</b>	tosca.artifacts.Implementation.Bash

#### 2689 5.4.4.3.1 Definition

```
tosca.artifacts.Implementation.Bash:  
  derived_from: tosca.artifacts.Implementation  
  description: Script artifact for the Unix Bash shell  
  mime_type: application/x-sh
```

```
file_ext: [ sh ]
```

#### 2690 **5.4.4.4 tosca.artifacts.Implementation.Python**

2691 This artifact type represents a Python file that contains Python language constructs that can be executed  
2692 within a Python interpreter.

<b>Shorthand Name</b>	Python
<b>Type Qualified Name</b>	tosca:Python
<b>Type URI</b>	tosca.artifacts.Implementation.Python

#### 2693 **5.4.4.4.1 Definition**

```
tosca.artifacts.Implementation.Python:  
  derived_from: tosca.artifacts.Implementation  
  description: Artifact for the interpreted Python language  
  mime_type: application/x-python  
  file_ext: [ py ]
```

### 2694 **5.5 Capabilities Types**

#### 2695 **5.5.1 tosca.capabilities.Root**

2696 This is the default (root) TOSCA Capability Type definition that all other TOSCA Capability Types derive  
2697 from.

#### 2698 **5.5.1.1 Definition**

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

#### 2699 **5.5.2 tosca.capabilities.Node**

2700 The Node capability indicates the base capabilities of a TOSCA Node Type.

<b>Shorthand Name</b>	Node
<b>Type Qualified Name</b>	tosca:Node
<b>Type URI</b>	tosca.capabilities.Node

#### 2701 **5.5.2.1 Definition**

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

2702 **5.5.3 tosca.capabilities.Compute**

2703 The Compute capability, when included on a Node Type or Template definition, indicates that the node  
2704 can provide hosting on a named compute resource.

<b>Shorthand Name</b>	Compute
<b>Type Qualified Name</b>	tosca:Compute
<b>Type URI</b>	tosca.capabilities.Compute

2705 **5.5.3.1 Properties**

Name	Required	Type	Constraints	Description
name	no	string	None	The optional name (or identifier) of a specific compute resource for hosting.
num_cpus	no	integer	greater_or_equal: 1	Number of (actual or virtual) CPUs associated with the Compute node.
cpu_frequency	no	scalar-unit.frequency	greater_or_equal: 0.1 GHz	Specifies the operating frequency of CPU's core. This property expresses the expected frequency of one (1) CPU as provided by the property "num_cpus".
disk_size	no	scalar-unit.size	greater_or_equal: 0 MB	Size of the local disk available to applications running on the Compute node (default unit is MB).
mem_size	no	scalar-unit.size	greater_or_equal: 0 MB	Size of memory available to applications running on the Compute node (default unit is MB).

2706 **5.5.3.2 Definition**

```
tosca.capabilities.Compute:  
  derived_from: tosca.capabilities.Root  
  properties:  
    name:  
      type: string  
      required: false  
    num_cpus:  
      type: integer  
      required: false  
      constraints:  
        - greater_or_equal: 1  
    cpu_frequency:  
      type: scalar-unit.frequency  
      required: false  
      constraints:  
        - greater_or_equal: 0.1 GHz  
    disk_size:
```

```

type: scalar-unit.size
required: false
constraints:
  - greater_or_equal: 0 MB
mem_size:
type: scalar-unit.size
required: false
constraints:
  - greater_or_equal: 0 MB

```

## 2707 5.5.4 [tosca.capabilities.Network](#)

2708 The Storage capability, when included on a Node Type or Template definition, indicates that the node can  
 2709 provide addressibility for the resource a named network with the specified ports.

<b>Shorthand Name</b>	Network
<b>Type Qualified Name</b>	tosca:Network
<b>Type URI</b>	tosca.capabilities.Network

### 2710 5.5.4.1 Properties

Name	Required	Type	Constraints	Description
name	no	string	None	The otional name (or identifier) of a specific network resource.

### 2711 5.5.4.2 Definition

```

tosca.capabilities.Network:
  derived_from: tosca.capabilities.Root
  properties:
    name:
      type: string
      required: false

```

## 2712 5.5.5 [tosca.capabilities.Storage](#)

2713 The Storage capability, when included on a Node Type or Template definition, indicates that the node can  
 2714 provide a named storage location with specified size range.

<b>Shorthand Name</b>	Storage
<b>Type Qualified Name</b>	tosca:Storage
<b>Type URI</b>	tosca.capabilities.Storage

2715 **5.5.5.1 Properties**

Name	Required	Type	Constraints	Description
name	no	string	None	The optional name (or identifier) of a specific storage resource.

2716 **5.5.5.2 Definition**

```
tosca.capabilities.Storage:
  derived_from: tosca.capabilities.Root
  properties:
    name:
      type: string
      required: false
```

2717 **5.5.6 tosca.capabilities.Compute.Container**

2718 The Container capability, when included on a Node Type or Template definition, indicates that the node  
 2719 can act as a container for (or a host for) one or more other declared Node Types.

<b>Shorthand Name</b>	Container
<b>Type Qualified Name</b>	tosca:Container tosca:Compute.Container
<b>Type URI</b>	tosca.capabilities.Container tosca.capabilities.Compute.Container

2720 **5.5.6.1 Properties**

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

2721 **5.5.6.2 Definition**

```
tosca.capabilities.compute.Container:
  derived_from: tosca.capabilities.Compute
```

2722 **5.5.7 tosca.capabilities.Endpoint**

2723 This is the default TOSCA type that should be used or extended to define a network endpoint capability.  
 2724 This includes the information to express a basic endpoint with a single port or a complex endpoint with

2725 multiple ports. By default the Endpoint is assumed to represent an address on a private network unless  
 2726 otherwise specified.

<b>Shorthand Name</b>	Endpoint
<b>Type Qualified Name</b>	tosca:Endpoint
<b>Type URI</b>	tosca.capabilities.Endpoint

2727 **5.5.7.1 Properties**

Name	Required	Type	Constraints	Description
protocol	yes	string	default: tcp	The name of the protocol (i.e., the protocol prefix) that the endpoint accepts (any OSI Layer 4-7 protocols)  Examples: http, https, ftp, tcp, udp, etc.
port	no	PortDef	greater_or_equal: 1 less_or_equal: 65535	The optional port of the endpoint.
secure	no	boolean	default: false	Requests for the endpoint to be secure and use credentials supplied on the ConnectsTo relationship.
url_path	no	string	None	The optional URL path of the endpoint's address if applicable for the protocol.
port_name	no	string	None	The optional name (or ID) of the network port this endpoint should be bound to.
network_name	no	string	default: PRIVATE	The optional name (or ID) of the network this endpoint should be bound to. network_name: PRIVATE   PUBLIC   <network_name>   <network_id>
initiator	no	string	one of: • source • target • peer  default: source	The optional indicator of the direction of the connection.
ports	no	map of PortSpec	None	The optional map of ports the Endpoint supports (if more than one)

2728 **5.5.7.2 Attributes**

Name	Required	Type	Constraints	Description
ip_address	yes	string	None	Note: This is the IP address as propagated up by the associated node's host (Compute) container.

2729 **5.5.7.3 Definition**

```
tosca.capabilities.Endpoint:
  derived_from: tosca.capabilities.Root
  properties:
```

```

protocol:
  type: string
  required: true
  default: tcp
port:
  type: PortDef
  required: false
secure:
  type: boolean
  required: false
  default: false
url_path:
  type: string
  required: false
port_name:
  type: string
  required: false
network_name:
  type: string
  required: false
  default: PRIVATE
initiator:
  type: string
  required: false
  default: source
  constraints:
    - valid_values: [ source, target, peer ]
ports:
  type: map
  required: false
  constraints:
    - min_length: 1
  entry_schema:
    type: PortSpec
attributes:
  ip_address:
    type: string

```

2730 **5.5.7.4 Additional requirements**

- 2731
- 2732
- Although both the port and ports properties are not required, one of port or ports must be provided in a valid [Endpoint](#).

2733 **5.5.8 tosca.capabilities.Endpoint.Public**

2734 This capability represents a public endpoint which is accessible to the general internet (and its public IP  
2735 address ranges).

2736 This public endpoint capability also can be used to create a floating (IP) address that the underlying  
2737 network assigns from a pool allocated from the application's underlying public network. This floating  
2738 address is managed by the underlying network such that can be routed an application's private address  
2739 and remains reliable to internet clients.

<b>Shorthand Name</b>	Endpoint.Public
<b>Type Qualified Name</b>	tosca:Endpoint.Public
<b>Type URI</b>	tosca.capabilities.Endpoint.Public

2740 **5.5.8.1 Definition**

```
tosca.capabilities.Endpoint.Public:  
  derived_from: tosca.capabilities.Endpoint  
  properties:  
    # Change the default network_name to use the first public network found  
    network_name:  
      type: string  
      default: PUBLIC  
      constraints:  
        - equal: PUBLIC  
    floating:  
      description: >  
        indicates that the public address should be allocated from a pool of  
        floating IPs that are associated with the network.  
      type: boolean  
      default: false  
      status: experimental  
    dns_name:  
      description: The optional name to register with DNS  
      type: string  
      required: false  
      status: experimental
```

2741 **5.5.8.2 Additional requirements**

- 2742
- 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.
- 2744
- 2745
- 2746

2747 **5.5.9 tosca.capabilities.Endpoint.Admin**

2748 This is the default TOSCA type that should be used or extended to define a specialized administrator  
2749 endpoint capability.

<b>Shorthand Name</b>	Endpoint.Admin
<b>Type Qualified Name</b>	tosca:Endpoint.Admin
<b>Type URI</b>	tosca.capabilities.Endpoint.Admin

2750 **5.5.9.1 Properties**

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

2751 **5.5.9.2 Definition**

```
tosca.capabilities.Endpoint.Admin:
  derived_from: tosca.capabilities.Endpoint
  # Change Endpoint secure indicator to true from its default of false
  properties:
    secure:
      type: boolean
      default: true
      constraints:
        - equal: true
```

2752 **5.5.9.3 Additional requirements**

- 2753 • TOSCA Orchestrator implementations of Endpoint.Admin (and connections to it) **SHALL** assure  
2754 that network-level security is enforced if possible.

2755 **5.5.10 tosca.capabilities.Endpoint.Database**

2756 This is the default TOSCA type that should be used or extended to define a specialized database  
2757 endpoint capability.

<b>Shorthand Name</b>	Endpoint.Database
<b>Type Qualified Name</b>	tosca:Endpoint.Database
<b>Type URI</b>	tosca.capabilities.Endpoint.Database

2758 **5.5.10.1 Properties**

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

2759 **5.5.10.2 Definition**

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

2760 **5.5.11 tosca.capabilities.Attachment**

2761 This is the default TOSCA type that should be used or extended to define an attachment capability of a  
 2762 (logical) infrastructure device node (e.g., [BlockStorage](#) node).

<b>Shorthand Name</b>	Attachment
<b>Type Qualified Name</b>	tosca:Attachment
<b>Type URI</b>	tosca.capabilities.Attachment

2763 **5.5.11.1 Properties**

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

2764 **5.5.11.2 Definition**

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

2765 **5.5.12 tosca.capabilities.OperatingSystem**

2766 This is the default TOSCA type that should be used to express an Operating System capability for a  
 2767 node.

<b>Shorthand Name</b>	OperatingSystem
<b>Type Qualified Name</b>	tosca:OperatingSystem
<b>Type URI</b>	tosca.capabilities.OperatingSystem

2768 **5.5.12.1 Properties**

Name	Required	Type	Constraints	Description
architecture	no	string	None	The Operating System (OS) architecture.  Examples of valid values include: x86_32, x86_64, etc.
type	no	string	None	The Operating System (OS) type.  Examples of valid values include: linux, aix, mac, windows, etc.
distribution	no	string	None	The Operating System (OS) distribution.  Examples of valid values for an "type" of "Linux" would include: debian, fedora, rhel and ubuntu.

Name	Required	Type	Constraints	Description
version	no	<a href="#">version</a>	None	The Operating System version.

2769 **5.5.12.2 Definition**

```

tosca.capabilities.OperatingSystem:
  derived_from: tosca.capabilities.Root
  properties:
    architecture:
      type: string
      required: false
    type:
      type: string
      required: false
    distribution:
      type: string
      required: false
    version:
      type: version
      required: false

```

2770 **5.5.12.3 Additional Requirements**

- 2771 • Please note that the string values for the properties **architecture**, **type** and **distribution**  
2772 SHALL be normalized to lowercase by processors of the service template for matching purposes.  
2773 For example, if a “**type**” value is set to either “Linux”, “LINUX” or “linux” in a service template, the  
2774 processor would normalize all three values to “linux” for matching purposes.

2775 **5.5.13 tosca.capabilities.Scalable**

2776 This is the default TOSCA type that should be used to express a scalability capability for a node.

<b>Shorthand Name</b>	Scalable
<b>Type Qualified Name</b>	tosca:Scalable
<b>Type URI</b>	tosca.capabilities.Scalable

2777 **5.5.13.1 Properties**

Name	Required	Type	Constraints	Description
min_instances	yes	<a href="#">integer</a>	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	<a href="#">integer</a>	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.

Name	Required	Type	Constraints	Description
default_instances	no	integer	N/A	<p>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.</p> <p><b>Note:</b> The value for this property MUST be in the range between the values set for 'min_instances' and 'max_instances' properties.</p>

2778 **5.5.13.2 Definition**

```

tosca.capabilities.Scalable:
  derived_from: tosca.capabilities.Root
  properties:
    min_instances:
      type: integer
      default: 1
    max_instances:
      type: integer
      default: 1
    default_instances:
      type: integer

```

2779 **5.5.13.3 Notes**

- 2780 • The actual number of instances for a node may be governed by a separate scaling policy which
- 2781 conceptually would be associated to either a scaling-capable node or a group of nodes in which it
- 2782 is defined to be a part of. This is a planned future feature of the TOSCA Simple Profile and not
- 2783 currently described.

2784 **5.5.14 tosca.capabilities.network.Bindable**

2785 A node type that includes the Bindable capability indicates that it can be bound to a logical network  
 2786 association via a network port.

<b>Shorthand Name</b>	network.Bindable
<b>Type Qualified Name</b>	tosca:network.Bindable
<b>Type URI</b>	tosca.capabilities.network.Bindable

2787 **5.5.14.1 Properties**

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

2788 **5.5.14.2 Definition**

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

2789 **5.6 Requirement Types**

2790 There are no normative Requirement Types currently defined in this working draft. Typically,  
2791 Requirements are described against a known Capability Type

2792 **5.7 Relationship Types**

2793 **5.7.1 tosca.relationships.Root**

2794 This is the default (root) TOSCA Relationship Type definition that all other TOSCA Relationship Types  
2795 derive from.

2796 **5.7.1.1 Attributes**

Name	Required	Type	Constraints	Description
tosca_id	yes	string	None	A unique identifier of the realized instance of a Relationship Template that derives from any TOSCA normative type.
tosca_name	yes	string	None	This attribute reflects the name of the Relationship Template as defined in the TOSCA service template. This name is not unique to the realized instance model of corresponding deployed application as each template in the model can result in one or more instances (e.g., scaled) when orchestrated to a provider environment.
state	yes	string	default: initial	The state of the relationship instance. See section "Relationship States" for allowed values.

2797 **5.7.1.2 Definition**

```
tosca.relationships.Root:
  description: The TOSCA root Relationship Type all other TOSCA base Relationship
  Types derive from
  attributes:
    tosca_id:
      type: string
    tosca_name:
      type: string
  interfaces:
    Configure:
      type: tosca.interfaces.relationship.Configure
```

2798 **5.7.2 tosca.relationships.DependsOn**

2799 This type represents a general dependency relationship between two nodes.

<b>Shorthand Name</b>	DependsOn
<b>Type Qualified Name</b>	tosca:DependsOn
<b>Type URI</b>	tosca.relationships.DependsOn

2800 **5.7.2.1 Definition**

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

2801 **5.7.3 tosca.relationships.HostedOn**

2802 This type represents a hosting relationship between two nodes.

<b>Shorthand Name</b>	HostedOn
<b>Type Qualified Name</b>	tosca:HostedOn
<b>Type URI</b>	tosca.relationships.HostedOn

2803 **5.7.3.1 Definition**

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

2804 **5.7.4 tosca.relationships.ConnectsTo**

2805 This type represents a network connection relationship between two nodes.

<b>Shorthand Name</b>	ConnectsTo
<b>Type Qualified Name</b>	tosca:ConnectsTo
<b>Type URI</b>	tosca.relationships.ConnectsTo

2806 **5.7.4.1 Definition**

```
tosca.relationships.ConnectsTo:
  derived_from: tosca.relationships.Root
  valid_target_types: [ tosca.capabilities.Endpoint ]
  properties:
    credential:
      type: tosca.datatypes.Credential
      required: false
```

2807 **5.7.4.2 Properties**

Name	Required	Type	Constraints	Description
credential	no	<a href="#">Credential</a>	None	The security credential to use to present to the target endpoint to for either authentication or authorization purposes.

2808 **5.7.5 tosca.relationships.AttachesTo**

2809 This type represents an attachment relationship between two nodes. For example, an AttachesTo  
 2810 relationship type would be used for attaching a storage node to a Compute node.

<b>Shorthand Name</b>	AttachesTo
<b>Type Qualified Name</b>	tosca:AttachesTo
<b>Type URI</b>	tosca.relationships.AttachesTo

2811 **5.7.5.1 Properties**

Name	Required	Type	Constraints	Description
location	yes	<a href="#">string</a>	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	<a href="#">string</a>	None	The logical device name which for the attached device (which is represented by the target node in the model). e.g., '/dev/hda1'

2812 **5.7.5.2 Attributes**

Name	Required	Type	Constraints	Description
device	no	<a href="#">string</a>	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.

2813 **5.7.5.3 Definition**

```
tosca.relationships.AttachesTo:
  derived_from: tosca.relationships.Root
  valid_target_types: [ tosca.capabilities.Attachment ]
  properties:
    location:
      type: string
      constraints:
        - min_length: 1
    device:
      type: string
```

```
required: false
```

## 2814 5.7.6 **tosca.relationships.RoutesTo**

2815 This type represents an intentional network routing between two Endpoints in different networks.

<b>Shorthand Name</b>	RoutesTo
<b>Type Qualified Name</b>	tosca:RoutesTo
<b>Type URI</b>	tosca.relationships.RoutesTo

### 2816 5.7.6.1 Definition

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

## 2817 5.8 Interface Types

2818 Interfaces are reusable entities that define a set of operations that that can be included as part of a Node  
2819 type or Relationship Type definition. Each named operations may have code or scripts associated with  
2820 them that orchestrators can execute for when transitioning an application to a given state.

### 2821 5.8.1 Additional Requirements

- 2822 • Designers of Node or Relationship types are not required to actually provide/associate code or  
2823 scripts with every operation for a given interface it supports. In these cases, orchestrators SHALL  
2824 consider that a “No Operation” or “no-op”.
- 2825 • The default behavior when providing scripts for an operation in a sub-type (sub-class) or a  
2826 template of an existing type which already has a script provided for that operation SHALL be  
2827 override. Meaning that the subclasses’ script is used in place of the parent type’s script.

### 2828 5.8.2 Best Practices

- 2829 • When TOSCA Orchestrators substitute an implementation for an abstract node in a deployed  
2830 service template it SHOULD be able to present a confirmation to the submitter to confirm the  
2831 implementation chosen would be acceptable.

### 2832 5.8.3 **tosca.interfaces.Root**

2833 This is the default (root) TOSCA Interface Type definition that all other TOSCA Interface Types derive  
2834 from.

#### 2835 5.8.3.1 Definition

```
tosca.interfaces.Root:  
  derived_from: tosca.entity.Root  
  description: The TOSCA root Interface Type all other TOSCA base Interface Types  
  derive from
```

2836 **5.8.4 tosca.interfaces.node.lifecycle.Standard**

2837 This lifecycle interface defines the essential, normative operations that TOSCA nodes may support.

<b>Shorthand Name</b>	Standard
<b>Type Qualified Name</b>	tosca: Standard
<b>Type URI</b>	tosca.interfaces.node.lifecycle.Standard

2838 **5.8.4.1 Definition**

```
tosca.interfaces.node.lifecycle.Standard:  
  derived_from: tosca.interfaces.Root  
  create:  
    description: Standard lifecycle create operation.  
  configure:  
    description: Standard lifecycle configure operation.  
  start:  
    description: Standard lifecycle start operation.  
  stop:  
    description: Standard lifecycle stop operation.  
  delete:  
    description: Standard lifecycle delete operation.
```

2839 **5.8.4.2 Create operation**

2840 The create operation is generally used to create the resource or service the node represents in the  
2841 topology. TOSCA orchestrators expect node templates to provide either a deployment artifact or an  
2842 implementation artifact of a defined artifact type that it is able to process. This specification defines  
2843 normative deployment and implementation artifact types all TOSCA Orchestrators are expected to be  
2844 able to process to support application portability.

2845 **5.8.4.3 TOSCA Orchestrator processing of Deployment artifacts**

2846 TOSCA Orchestrators, when encountering a deployment artifact on the create operation; will  
2847 automatically attempt to deploy the artifact based upon its artifact type. This means that no  
2848 implementation artifacts (e.g., scripts) are needed on the create operation to provide commands that  
2849 deploy or install the software.

2850  
2851 For example, if a TOSCA Orchestrator is processing an application with a node of type  
2852 SoftwareComponent and finds that the node's template has a create operation that provides a filename  
2853 (or references to an artifact which describes a file) of a known TOSCA deployment artifact type such as  
2854 an Open Virtualization Format (OVF) image it will automatically deploy that image into the  
2855 SoftwareComponent's host Compute node.

2856 **5.8.4.4 Operation sequencing and node state**


2857 The following diagrams show how TOSCA orchestrators sequence the operations of the Standard  
2858 lifecycle in normal node startup and shutdown procedures.

2859 The following key should be used to interpret the diagrams:

Operation Invocation	
Node State	
Transition State	


2860 **5.8.4.4.1 Normal node startup sequence diagram**

2861 The following diagram shows how the TOSCA orchestrator would invoke operations on the Standard  
 2862 lifecycle to startup a node.



2863 **5.8.4.4.2 Normal node shutdown sequence diagram**

2864 The following diagram shows how the TOSCA orchestrator would invoke operations on the Standard  
 2865 lifecycle to shut down a node.



2866

### 2867 5.8.5 [tosca.interfaces.relationship.Configure](#)

2868 The lifecycle interfaces define the essential, normative operations that each TOSCA Relationship Types  
 2869 may support.

<b>Shorthand Name</b>	Configure
<b>Type Qualified Name</b>	tosca:Configure
<b>Type URI</b>	tosca.interfaces.relationship.Configure

#### 2870 5.8.5.1 Definition


```

tosca.interfaces.relationship.Configure:
  derived_from: tosca.interfaces.Root
  pre_configure_source:
    description: Operation to pre-configure the source endpoint.
  pre_configure_target:
    description: Operation to pre-configure the target endpoint.
  post_configure_source:
    description: Operation to post-configure the source endpoint.
  post_configure_target:
    description: Operation to post-configure the target endpoint.
  add_target:
    description: Operation to notify the source node of a target node being added
    via a relationship.
  add_source:
    description: Operation to notify the target node of a source node which is
    now available via a relationship.
  description:
  target_changed:
    description: Operation to notify source some property or attribute of the
    target changed
  
```

```
remove_target:  
  description: Operation to remove a target node.
```

2871


2872 **5.8.5.2 Invocation Conventions**



2873 TOSCA relationships are directional connecting a source node to a target node. When TOSCA  
 2874 Orchestrator connects a source and target node together using a relationship that supports the Configure  
 2875 interface it will “interleave” the operations invocations of the Configure interface with those of the node’s  
 2876 own Standard lifecycle interface. This concept is illustrated below:

2877 **5.8.5.3 Normal node start sequence with Configure relationship operations**

2878 The following diagram shows how the TOSCA orchestrator would invoke Configure lifecycle operations in  
 2879 conjunction with Standard lifecycle operations during a typical startup sequence on a node.



#### 2880 **5.8.5.4 Node-Relationship configuration sequence**

2881 Depending on which side (i.e., source or target) of a relationship a node is on, the orchestrator will:

- 2882 • Invoke either the **pre\_configure\_source** or **pre\_configure\_target** operation as supplied by
- 2883 the relationship on the node.
- 2884 • Invoke the node's **configure** operation.
- 2885 • Invoke either the **post\_configure\_source** or **post\_configure\_target** as supplied by the
- 2886 relationship on the node.

2887 Note that the **pre\_configure\_xxx** and **post\_configure\_xxx** are invoked only once per node instance.

#### 2888 **5.8.5.4.1 Node-Relationship add, remove and changed sequence**

2889 Since a topology template contains nodes that can dynamically be added (and scaled), removed or  
2890 changed as part of an application instance, the Configure lifecycle includes operations that are invoked  
2891 on node instances that to notify and address these dynamic changes.

2892

2893 For example, a source node, of a relationship that uses the Configure lifecycle, will have the relationship  
2894 operations **add\_target**, or **remove\_target** invoked on it whenever a target node instance is added or  
2895 removed to the running application instance. In addition, whenever the node state of its target node  
2896 changes, the **target\_changed** operation is invoked on it to address this change. Conversely, the  
2897 **add\_source** and **remove\_source** operations are invoked on the source node of the relationship.

#### 2898 **5.8.5.5 Notes**

- 2899 • The target (provider) MUST be active and running (i.e., all its dependency stack MUST be
- 2900 fulfilled) prior to invoking **add\_target**
- 2901 • In other words, all Requirements MUST be satisfied before it advertises its capabilities (i.e.,
- 2902 the attributes of the matched Capabilities are available).
- 2903 • In other words, it cannot be “consumed” by any dependent node.
- 2904 • Conversely, since the source (consumer) needs information (attributes) about any targets
- 2905 (and their attributes) being removed before it actually goes away.
- 2906 • The **remove\_target** operation should only be executed if the target has had **add\_target**
- 2907 executed. BUT in truth we're first informed about a target in **pre\_configure\_source**, so if we
- 2908 execute that the source node should see **remove\_target** called to cleanup.
- 2909 • **Error handling:** If any node operation of the topology fails processing should stop on that node
- 2910 template and the failing operation (script) should return an error (failure) code when possible.

### 2911 **5.9 Node Types**

#### 2912 **5.9.1 tosca.nodes.Root**

2913 The TOSCA **Root** Node Type is the default type that all other TOSCA base Node Types derive from.  
2914 This allows for all TOSCA nodes to have a consistent set of features for modeling and management (e.g.,  
2915 consistent definitions for requirements, capabilities and lifecycle interfaces).

2916

<b>Shorthand Name</b>	Root
<b>Type Qualified Name</b>	tosca:Root
<b>Type URI</b>	tosca.nodes.Root

2917 **5.9.1.1 Properties**

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

2918 **5.9.1.2 Attributes**

Name	Required	Type	Constraints	Description
tosca_id	yes	string	None	A unique identifier of the realized instance of a Node Template that derives from any TOSCA normative type.
tosca_name	yes	string	None	This attribute reflects the name of the Node Template as defined in the TOSCA service template. This name is not unique to the realized instance model of corresponding deployed application as each template in the model can result in one or more instances (e.g., scaled) when orchestrated to a provider environment.
state	yes	string	default: initial	The state of the node instance. See section “ <a href="#">Node States</a> ” for allowed values.

2919 **5.9.1.3 Definition**

```

tosca.nodes.Root:
  derived_from: toasca.entity.Root
  description: The TOSCA Node Type all other TOSCA base Node Types derive from
  attributes:
    toasca_id:
      type: string
    toasca_name:
      type: string
    state:
      type: string
  capabilities:
    feature:
      type: tosca.capabilities.Node
  requirements:
    - dependency:
      capability: tosca.capabilities.Node
      node: tosca.nodes.Root
      relationship: tosca.relationships.DependsOn
      occurrences: [ 0, UNBOUNDED ]

```

```

interfaces:
  Standard:
    type: tosca.interfaces.node.lifecycle.Standard

```

2920 **5.9.1.4 Additional Requirements**

- 2921 • All Node Type definitions that wish to adhere to the TOSCA Simple Profile **SHOULD** extend from the  
 2922 TOSCA Root Node Type to be assured of compatibility and portability across implementations.

2923 **5.9.2 tosca.nodes.Compute**

2924 The TOSCA **Compute** node represents one or more real or virtual processors of software applications or  
 2925 services along with other essential local resources. Collectively, the resources the compute node  
 2926 represents can logically be viewed as a (real or virtual) “server”.

<b>Shorthand Name</b>	Compute
<b>Type Qualified Name</b>	tosca:Compute
<b>Type URI</b>	tosca.nodes.Compute

2927 **5.9.2.1 Properties**

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

2928 **5.9.2.2 Attributes**

Name	Required	Type	Constraints	Description
private_address	no	<a href="#">string</a>	None	The primary private IP address assigned by the cloud provider that applications may use to access the Compute node.
public_address	no	<a href="#">string</a>	None	The primary public IP address assigned by the cloud provider that applications may use to access the Compute node.
networks	no	<a href="#">map of NetworkInfo</a>	None	The list of logical networks assigned to the compute host instance and information about them.
ports	no	<a href="#">map of PortInfo</a>	None	The list of logical ports assigned to the compute host instance and information about them.

2929 **5.9.2.3 Definition**

```

tosca.nodes.Compute:
  derived_from: tosca.nodes.Root
  attributes:
    private_address:
      type: string
    public_address:

```

```

    type: string
  networks:
    type: map
    entry_schema:
      type: tosca.datatypes.network.NetworkInfo
  ports:
    type: map
    entry_schema:
      type: tosca.datatypes.network.PortInfo
  requirements:
    - local_storage:
        capability: tosca.capabilities.Attachment
        node: tosca.nodes.BlockStorage
        relationship: tosca.relationships.AttachesTo
        occurrences: [0, UNBOUNDED]
  capabilities:
    host:
      type: tosca.capabilities.Container
      valid_source_types: [tosca.nodes.SoftwareComponent]
    endpoint:
      type: tosca.capabilities.Endpoint.Admin
    os:
      type: tosca.capabilities.OperatingSystem
    scalable:
      type: tosca.capabilities.Scalable
    binding:
      type: tosca.capabilities.network.Bindable

```

#### 2930 **5.9.2.4 Additional Requirements**

- 2931
- The underlying implementation of the Compute node SHOULD have the ability to instantiate
- 2932 guest operating systems (either actual or virtualized) based upon the OperatingSystem capability
- 2933 properties if they are supplied in the a node template derived from the Compute node type.

#### 2934 **5.9.3 [tosca.nodes.SoftwareComponent](#)**

2935 The TOSCA **SoftwareComponent** node represents a generic software component that can be managed

2936 and run by a TOSCA **Compute Node Type**.

<b>Shorthand Name</b>	SoftwareComponent
<b>Type Qualified Name</b>	tosca:SoftwareComponent
<b>Type URI</b>	tosca.nodes.SoftwareComponent

2937 **5.9.3.1 Properties**

Name	Required	Type	Constraints	Description
component_version	no	<a href="#">version</a>	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.

2938 **5.9.3.2 Attributes**

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

2939 **5.9.3.3 Definition**

```

tosca.nodes.SoftwareComponent:
  derived_from: tosca.nodes.Root
  properties:
    # domain-specific software component version
    component_version:
      type: version
      required: false
    admin_credential:
      type: tosca.datatypes.Credential
      required: false
  requirements:
    - host:
      capability: tosca.capabilities.Container
      node: tosca.nodes.Compute
      relationship: tosca.relationships.HostedOn

```

2940 **5.9.3.4 Additional Requirements**

- 2941
- Nodes that can directly be managed and run by a TOSCA **Compute** Node Type **SHOULD** extend from this type.
- 2942

2943 **5.9.4 tosca.nodes.WebServer**

2944 This TOSA **WebServer** Node Type represents an abstract software component or service that is capable  
2945 of hosting and providing management operations for one or more **WebApplication** nodes.

<b>Shorthand Name</b>	WebServer
<b>Type Qualified Name</b>	tosca:WebServer
<b>Type URI</b>	tosca.nodes.WebServer

2946 **5.9.4.1 Properties**

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

2947 **5.9.4.2 Definition**

```
tosca.nodes.WebServer:
  derived_from: tosca.nodes.SoftwareComponent
  capabilities:
    # Private, layer 4 endpoints
    data_endpoint: tosca.capabilities.Endpoint
    admin_endpoint: tosca.capabilities.Endpoint.Admin
  host:
    type: tosca.capabilities.Container
    valid_source_types: [ tosca.nodes.WebApplication ]
```

2948 **5.9.4.3 Additional Requirements**

- 2949 • This node **SHALL** export both a secure endpoint capability (i.e., **admin\_endpoint**), typically for  
2950 administration, as well as a regular endpoint (i.e., **data\_endpoint**) for serving data.

2951 **5.9.5 tosca.nodes.WebApplication**

2952 The TOSCA **WebApplication** node represents a software application that can be managed and run by a  
2953 TOSCA **WebServer** node. Specific types of web applications such as Java, etc. could be derived from  
2954 this type.

<b>Shorthand Name</b>	WebApplication
<b>Type Qualified Name</b>	tosca: WebApplication
<b>Type URI</b>	tosca.nodes.WebApplication

2955 **5.9.5.1 Properties**

Name	Required	Type	Constraints	Description
context_root	no	<a href="#">string</a>	None	The web application's context root which designates the application's URL path within the web server it is hosted on.

2956 **5.9.5.2 Definition**

```

tosca.nodes.WebApplication:
  derived_from: tosca.nodes.Root
  properties:
    context_root:
      type: string
  capabilities:
    app_endpoint:
      type: tosca.capabilities.Endpoint
  requirements:
    - host:
      capability: tosca.capabilities.Container
      node: tosca.nodes.WebServer
      relationship: tosca.relationships.HostedOn

```

2957 **5.9.6 tosca.nodes.DBMS**

2958 The TOSCA **DBMS** node represents a typical relational, SQL Database Management System software  
2959 component or service.

2960 **5.9.6.1 Properties**

Name	Required	Type	Constraints	Description
root_password	no	<a href="#">string</a>	None	The optional root password for the DBMS server.
port	no	<a href="#">integer</a>	None	The DBMS server's port.

2961 **5.9.6.2 Definition**

```

tosca.nodes.DBMS:
  derived_from: tosca.nodes.SoftwareComponent
  properties:
    root_password:
      type: string
      required: false
      description: the optional root password for the DBMS service
    port:
      type: integer
      required: false
      description: the port the DBMS service will listen to for data and requests
  capabilities:
    host:
      type: tosca.capabilities.Container
  valid_source_types: [ tosca.nodes.Database ]

```

2962 **5.9.7 tosca.nodes.Database**

2963 The TOSCA Database node represents a logical database that can be managed and hosted by a TOSCA  
2964 DBMS node.

<b>Shorthand Name</b>	Database
<b>Type Qualified Name</b>	tosca:Database
<b>Type URI</b>	tosca.nodes.Database

2965 **5.9.7.1 Properties**

Name	Required	Type	Constraints	Description
name	yes	string	None	The logical database Name
port	no	integer	None	The port the database service will use to listen for incoming data and requests.
user	no	string	None	The special user account used for database administration.
password	no	string	None	The password associated with the user account provided in the 'user' property.

2966 **5.9.7.2 Definition**

```

tosca.nodes.Database:
  derived_from: tosca.nodes.Root
  properties:
    name:
      type: string
      description: the logical name of the database
    port:
      type: integer
      description: the port the underlying database service will listen to for
data
    user:
      type: string
      description: the optional user account name for DB administration
      required: false
    password:
      type: string
      description: the optional password for the DB user account
      required: false
  requirements:
    - host:
      capability: tosca.capabilities.Container
      node: tosca.nodes.DBMS
      relationship: tosca.relationships.HostedOn
  capabilities:
    database_endpoint:
      type: tosca.capabilities.Endpoint.Database

```

2967 **5.9.8 tosca.nodes.Storage.ObjectStorage**

2968 The TOSCA **ObjectStorage** node represents storage that provides the ability to store data as objects (or  
2969 BLOBs of data) without consideration for the underlying filesystem or devices.

<b>Shorthand Name</b>	ObjectStorage
<b>Type Qualified Name</b>	tosca:ObjectStorage
<b>Type URI</b>	tosca.nodes.Storage.ObjectStorage

2970 **5.9.8.1 Properties**

Name	Required	Type	Constraints	Description
name	yes	<a href="#">string</a>	None	The logical name of the object store (or container).

Name	Required	Type	Constraints	Description
size	no	<a href="#">scalar-unit.size</a>	greater_or_equal: 0 GB	The requested initial storage size (default unit is in Gigabytes).
maxsize	no	<a href="#">scalar-unit.size</a>	greater_or_equal: 0 GB	The requested maximum storage size (default unit is in Gigabytes).

2971 **5.9.8.2 Definition**

```

tosca.nodes.Storage.ObjectStorage:
  derived_from: tosca.nodes.Root
  properties:
    name:
      type: string
    size:
      type: scalar-unit.size
      constraints:
        - greater_or_equal: 0 GB
    maxsize:
      type: scalar-unit.size
      constraints:
        - greater_or_equal: 0 GB
  capabilities:
    storage_endpoint:
      type: tosca.capabilities.Endpoint

```

2972 **5.9.8.3 Notes:**

- 2973 • Subclasses of the `tosca.nodes.ObjectStorage` node type may impose further constraints on  
2974 properties. For example, a subclass may constrain the (minimum or maximum) length of the  
2975 'name' property or include a regular expression to constrain allowed characters used in the  
2976 'name' property.

2977 **5.9.9 `tosca.nodes.Storage.BlockStorage`**

2978 The TOSCA **BlockStorage** node currently represents a server-local block storage device (i.e., not  
2979 shared) offering evenly sized blocks of data from which raw storage volumes can be created.

2980 **Note:** In this draft of the TOSCA Simple Profile, distributed or Network Attached Storage (NAS) are not  
2981 yet considered (nor are clustered file systems), but the TC plans to do so in future drafts.

<b>Shorthand Name</b>	BlockStorage
<b>Type Qualified Name</b>	tosca:BlockStorage
<b>Type URI</b>	tosca.nodes.Storage.BlockStorage

2982 **5.9.9.1 Properties**

Name	Required	Type	Constraints	Description
size	yes *	scalar-unit.size	greater_or_equal: 1 MB	The requested storage size (default unit is MB). * <b>Note:</b> <ul style="list-style-type: none"> <li>• <b>Required</b> when an existing volume (i.e., volume_id) is not available.</li> <li>• If <b>volume_id</b> is provided, size is ignored. Resize of existing volumes is not considered at this time.</li> </ul>
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).

2983 **5.9.9.2 Attributes**

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

2984 **5.9.9.3 Definition**

```
tosca.nodes.Storage.BlockStorage:
  derived_from: tosca.nodes.Root
  properties:
    size:
      type: scalar-unit.size
      constraints:
        - greater_or_equal: 1 MB
    volume_id:
      type: string
      required: false
    snapshot_id:
      type: string
      required: false
  capabilities:
    attachment:
      type: tosca.capabilities.Attachment
```

2985 **5.9.9.4 Additional Requirements**

- 2986
- The **size** property is required when an existing volume (i.e., **volume\_id**) is not available.
- 2987 However, if the property **volume\_id** is provided, the **size** property is ignored.

2988 **5.9.9.5 Notes**

- 2989
- Resize is of existing volumes is not considered at this time.
- 2990
- 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).
- 2991
- Currently, this version of the Simple Profile does not consider regions (or availability zones) when modeling storage.
- 2992
- 2993
- 2994

2995 **5.9.10 tosca.nodes.Container.Runtime**

2996 The TOSCA **Container** Runtime node represents operating system-level virtualization technology used  
2997 to run multiple application services on a single Compute host.

<b>Shorthand Name</b>	Container.Runtime
<b>Type Qualified Name</b>	tosca:Container.Runtime
<b>Type URI</b>	tosca.nodes.Container.Runtime

2998 **5.9.10.1 Definition**

```
tosca.nodes.Container.Runtime:  
  derived_from: tosca.nodes.SoftwareComponent  
  capabilities:  
    host:  
      type: tosca.capabilities.Container  
    scalable:  
      type: tosca.capabilities.Scalable
```

2999 **5.9.11 tosca.nodes.Container.Application**

3000 The TOSCA **Container** Application node represents an application that requires **Container**-level  
3001 virtualization technology.

<b>Shorthand Name</b>	Container.Application
<b>Type Qualified Name</b>	tosca:Container.Application
<b>Type URI</b>	tosca.nodes.Container.Application

3002 **5.9.11.1 Definition**

```
tosca.nodes.Container.Application:  
  derived_from: tosca.nodes.Root  
  requirements:
```

- host:
  - capability: [tosca.capabilities.Container](#)
  - node: [tosca.nodes.Container.Runtime](#)
  - relationship: [tosca.relationships.HostedOn](#)
- storage:
  - capability: [tosca.capabilities.Storage](#)
- network:
  - capability: [tosca.capabilities.EndPoint](#)

## 3003 5.9.12 [tosca.nodes.LoadBalancer](#)

3004 The TOSCA **Load Balancer** node represents logical function that be used in conjunction with a Floating  
 3005 Address to distribute an application's traffic (load) across a number of instances of the application (e.g.,  
 3006 for a clustered or scaled application).

<b>Shorthand Name</b>	LoadBalancer
<b>Type Qualified Name</b>	tosca:LoadBalancer
<b>Type URI</b>	tosca.nodes.LoadBalancer

### 3007 5.9.12.1 Definition

```
tosca.nodes.LoadBalancer:
  derived_from: tosca.nodes.Root
  properties:
    algorithm:
      type: string
      required: false
      status: experimental
  capabilities:
    client:
      type: tosca.capabilities.Endpoint.Public
      occurrences: [0, UNBOUNDED]
      description: the Floating (IP) client's on the public network can connect
to
  requirements:
    - application:
      capability: tosca.capabilities.Endpoint
      relationship: tosca.relationships.RoutesTo
      occurrences: [0, UNBOUNDED]
      description: Connection to one or more load balanced applications
```

3008 **5.9.12.2 Notes:**

- 3009
- 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.
- 3010
- 3011

3012 **5.10 Group Types**

3013 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.

3014

3015

3016

3017

3018 **Note:** Additional normative TOSCA Group Types and use cases for them will be developed in future drafts of this specification.

3019

3020 **5.10.1 `tosca.groups.Root`**

3021 This is the default (root) TOSCA **Group Type** definition that all other TOSCA base Group Types derive from.

3022

3023 **5.10.1.1 Definition**

```
tosca.groups.Root:
  description: The TOSCA Group Type all other TOSCA Group Types derive from
  interfaces:
  Standard:
    type: tosca.interfaces.node.lifecycle.Standard
```

3024 **5.10.1.2 Notes:**

- 3025
- 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.
- 3026
- 3027

3028 **5.11 Policy Types**

3029 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.

3030

3031

3032

3033 **5.11.1 `tosca.policies.Root`**

3034 This is the default (root) TOSCA Policy Type definition that all other TOSCA base Policy Types derive from.

3035

3036 **5.11.1.1 Definition**

```
tosca.policies.Root:
  description: The TOSCA Policy Type all other TOSCA Policy Types derive from
```

3037 **5.11.2 tosca.policies.Placement**

3038 This is the default (root) TOSCA Policy Type definition that is used to govern placement of TOSCA nodes  
3039 or groups of nodes.

3040 **5.11.2.1 Definition**

```
tosca.policies.Placement:  
  derived_from: tosca.policies.Root  
  description: The TOSCA Policy Type definition that is used to govern placement  
of TOSCA nodes or groups of nodes.
```

3041 **5.11.3 tosca.policies.Scaling**

3042 This is the default (root) TOSCA Policy Type definition that is used to govern scaling of TOSCA nodes or  
3043 groups of nodes.

3044 **5.11.3.1 Definition**

```
tosca.policies.Scaling:  
  derived_from: tosca.policies.Root  
  description: The TOSCA Policy Type definition that is used to govern scaling of  
TOSCA nodes or groups of nodes.
```

3045 **5.11.4 tosca.policies.Update**

3046 This is the default (root) TOSCA Policy Type definition that is used to govern update of TOSCA nodes or  
3047 groups of nodes.

3048 **5.11.4.1 Definition**

```
tosca.policies.Update:  
  derived_from: tosca.policies.Root  
  description: The TOSCA Policy Type definition that is used to govern update of  
TOSCA nodes or groups of nodes.
```

3049 **5.11.5 tosca.policies.Performance**

3050 This is the default (root) TOSCA Policy Type definition that is used to declare performance requirements  
3051 for TOSCA nodes or groups of nodes.

3052 **5.11.5.1 Definition**

```
tosca.policies.Performance:  
  derived_from: tosca.policies.Root  
  description: The TOSCA Policy Type definition that is used to declare  
performance requirements for TOSCA nodes or groups of nodes.
```

3053

3054

## 6 TOSCA Cloud Service Archive (CSAR) format

3055  
3056

Except for the examples, this section is **normative** and defines changes to the TOSCA archive format relative to the TOSCA v1.0 XML specification.

3057

3058  
3059  
3060  
3061  
3062

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.

3063

### 6.1 Overall Structure of a CSAR

3064

A CSAR zip file is required to contain one of the following:

3065  
3066  
3067  
3068  
3069

- 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.
- a yaml (.yaml or .yml) file at the root of the archive. The yaml file being a valid tosca definition template that **MUST** define a metadata section where `template_name` and `template_version` are required.

3070  
3071  
3072

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.

3073

### 6.2 TOSCA Meta File

3074  
3075  
3076  
3077

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.

3078  
3079  
3080  
3081

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).

3082  
3083  
3084  
3085

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.

3086  
3087

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**.

3088

#### 6.2.1 Example

3089  
3090

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
```

3091

3092 This **TOSCA.meta** file indicates its simplified TOSCA Simple Profile structure by means of the **CSAR-**  
3093 **Version** keyword with value **1.1**. The **Entry-Definitions** keyword points to a TOSCA definitions  
3094 YAML file with the name **tosca\_elk.yaml** which is contained in a directory called **definitions** within  
3095 the root of the CSAR file.

## 3096 **6.3 Archive without TOSCA-Metadata**

3097 In case the archive doesn't contains a TOSCA-Metadata directory the archive is required to contains a  
3098 single YAML file at the root of the archive (other templates may exists in sub-directories).

3099 This file must be a valid TOSCA definitions YAML file with the additional restriction that the metadata  
3100 section (as defined in 3.9.3.2) is required and **template\_name** and **template\_version** metadata are also  
3101 required.

3102 TOSCA processors should recognized this file as being the CSAR Entry-Definitions file. The CSAR-  
3103 Version is defined by the **template\_version** metadata section. The **Created-By** value is defined by the  
3104 **template\_author** metadata.

### 3105 **6.3.1 Example**

3106 The following represents a valid TOSCA template file acting as the CSAR Entry-Definitions file in an  
3107 archive without TOSCA-Metadata directory.

```
tosca_definitions_version: tosca_simple_yaml_1_1

metadata:
  template_name: my_template
  template_author: OASIS TOSCA TC
  template_version: 1.0
```

3108

3109

## 7 TOSCA workflows

3110 TOSCA defines two different kinds of workflows that can be used to deploy (instantiate and start),  
3111 manage at runtime or undeploy (stop and delete) a TOSCA topology: declarative workflows and  
3112 imperative workflows. Declarative workflows are automatically generated by the TOSCA orchestrator  
3113 based on the nodes, relationships, and groups defined in the topology. Imperative workflows are manually  
3114 specified by the author of the topology and allows the specification of any use-case that has not been  
3115 planned in the definition of node and relationships types or for advanced use-case (including reuse of  
3116 existing scripts and workflows).

3117

3118 Workflows can be triggered on deployment of a topology (deploy workflow) on undeployment (undeploy  
3119 workflow) or during runtime, manually, or automatically based on policies defined for the topology.

3120

3121 **Note:** The TOSCA orchestrators will execute a single workflow at a time on a topology to guarantee that  
3122 the defined workflow can be consistent and behave as expected.

### 7.1 Normative workflows

3124 TOSCA defines several normative workflows that are used to operate a Topology. That is, reserved  
3125 names of workflows that should be preserved by TOSCA orchestrators and that, if specified in the  
3126 topology will override the workflow generated by the orchestrator :

- 3127 • **deploy**: is the workflow used to instantiate and perform the initial deployment of the topology.
- 3128 • **undeploy**: is the workflow used to remove all instances of a topology.

#### 7.1.1 Notes

3130 Future versions of the specification will describe the normative naming and declarative generation of  
3131 additional workflows used to operate the topology at runtime.

- 3132 • **scaling workflows**: defined for every scalable nodes or based on scaling policies
- 3133 • **auto-healing workflows**: defined in order to restart nodes that may have failed

### 7.2 Declarative workflows

3135 Declarative workflows are the result of the weaving of topology's node, relationships, and groups  
3136 workflows.

3137 The weaving process generates the workflow of every single node in the topology, insert operations from  
3138 the relationships and groups and finally add ordering consideration. The weaving process will also take  
3139 care of the specific lifecycle of some nodes and the TOSCA orchestrator is responsible to trigger errors or  
3140 warnings in case the weaving cannot be processed or lead to cycles for example.

3141 This section aims to describe and explain how a TOSCA orchestrator will generate a workflow based on  
3142 the topology entities (nodes, relationships and groups).

#### 7.2.1 Notes

3144 This section details specific constraints and considerations that applies during the weaving process.

##### 7.2.1.1 Orchestrator provided nodes lifecycle and weaving

3146 When a node is abstract the orchestrator is responsible for providing a valid matching resources for the  
3147 node in order to deploy the topology. This consideration is also valid for dangling requirements (as they  
3148 represents a quick way to define an actual node).

3149 The lifecycle of such nodes is the responsibility of the orchestrator and they may not answer to the  
3150 normative TOSCA lifecycle. Their workflow is considered as "delegate" and acts as a black-box between  
3151 the initial and started state in the install workflow and the started to deleted states in the uninstall  
3152 workflow.

3153 If a relationship to some of this node defines operations or lifecycle dependency constraint that relies on  
3154 intermediate states, the weaving SHOULD fail and the orchestrator SHOULD raise an error.

## 3155 7.2.2 Relationship impacts on topology weaving

3156 This section explains how relationships impacts the workflow generation to enable the composition of  
3157 complex topologies.

### 3158 7.2.2.1 toska.relationships.DependsOn


3159 The depends on relationship is used to establish a dependency from a node to another. A source node  
3160 that depends on a target node will be created only after the other entity has been started.

#### 3161 7.2.2.2 Note

3162 DependsOn relationship SHOULD not be implemented. Even if the Configure interface can be  
3163 implemented this is not considered as a best-practice. If you need specific implementation, please have a  
3164 look at the ConnectsTo relationship.

#### 3165 7.2.2.2.1 Example DependsOn


3166 This example show the usage of a generic DependsOn relationship between two custom software  
3167 components.  
3168



3169 In this example the relationship configure interface doesn't define operations so they don't appear in the  
3170 generated lifecycle.  
3171

3172 **7.2.2.3 tosca.relationships.ConnectsTo**

3173 The connects to relationship is similar to the DependsOn relationship except that it is intended to provide  
 3174 an implementation. The difference is more theoretical than practical but helps users to make an actual  
 3175 distinction from a meaning perspective.



3176

3177 **7.2.2.4 tosca.relationships.HostedOn**

3178 The hosted\_on dependency relationship allows to define a hosting relationship between an entity and  
 3179 another. The hosting relationship has multiple impacts on the workflow and execution:


- 3180 • The implementation artifacts of the source node is executed on the same host as the one of the
- 3181 target node.
- 3182 • The create operation of the source node is executed only once the target node reach the started
- 3183 state.
- 3184 • When multiple nodes are hosted on the same host node, the defined operations will not be
- 3185 executed concurrently even if the theoretical workflow could allow it (actual generated workflow
- 3186 will avoid concurrency).

3187 **7.2.2.4.1 Example Software Component HostedOn Compute**

3188 This example explain the TOSCA weaving operation of a custom SoftwareComponent on a  
 3189 tosca.nodes.Compute instance. The compute node is an orchestrator provided node meaning that it's  
 3190 lifecycle is delegated to the orchestrator. This is a black-box and we just expect a started compute node  
 3191 to be provided by the orchestrator.

3192 The software node lifecycle operations will be executed on the Compute node (host) instance.

3193




3194

### 3195 7.2.2.4.2 Example Software Component HostedOn Software Component

3196 Tosca allows some more complex hosting scenarios where a software component could be hosted on  
3197 another software component.

3198



3199

3200 In such scenarios the software create operation is triggered only once the software\_base node has  
3201 reached the started state.

### 3202 7.2.2.4.3 Example 2 Software Components HostedOn Compute

3203 This example illustrate concurrency constraint introduced by the management of multiple nodes on a  
3204 single compute.

## 3205 7.2.3 Limitations

### 3206 7.2.3.1 Hosted nodes concurrency

3207 TOSCA implementation currently does not allow concurrent executions of scripts implementation artifacts  
3208 (shell, python, ansible, puppet, chef etc.) on a given host. This limitation is not applied on multiple hosts.  
3209 This limitation is expressed through the HostedOn relationship limitation expressing that when multiple  
3210 components are hosted on a given host node then their operations will not be performed concurrently  
3211 (generated workflow will ensure that operations are not concurrent).

### 3212 7.2.3.2 Dependent nodes concurrency

3213 When a node depends on another node no operations will be processed concurrently. In some situations,  
3214 especially when the two nodes lies on different hosts we could expect the create operation to be executed  
3215 concurrently for performance optimization purpose. The current version of the specification will allow to  
3216 use imperative workflows to solve this use-case. However, this scenario is one of the scenario that we  
3217 want to improve and handle in the future through declarative workflows.

### 3218 7.2.3.3 Target operations and get\_attribute on source

3219 The current ConnectsTo workflow implies that the target node is started before the source node is even  
3220 created. This means that pre\_configure\_target and post\_configure\_target operations cannot use any  
3221 input based on source attribute. It is however possible to refer to get\_property inputs based on source  
3222 properties. For advanced configurations the add\_source operation should be used.

3223 Note also that future plans on declarative workflows improvements aims to solve this kind of issues while  
3224 it is currently possible to use imperative workflows.

## 3225 7.3 Imperative workflows

3226 Imperative workflows are user defined and can define any really specific constraints and ordering of  
3227 activities. They are really flexible and powerful and can be used for any complex use-case that cannot be  
3228 solved in declarative workflows. However, they provide less reusability as they are defined for a specific  
3229 topology rather than being dynamically generated based on the topology content.

### 3230 7.3.1 Defining sequence of operations in an imperative workflow

3231 Imperative workflow grammar defines two ways to define the sequence of operations in an imperative  
3232 workflow:

- 3233 • Leverage the **on\_success** definition to define the next steps that will be executed in parallel.
- 3234 • Leverage a sequence of activity in a step.

#### 3235 7.3.1.1 Using on\_success to define steps ordering

3236 The graph of workflow steps is build based on the values of **on\_success** elements of the various defined  
3237 steps. The graph is built based on the following rules:

- 3238 • All steps that defines an **on\_success** operation must be executed before the next step can be  
3239 executed. So if A and C defines an **on\_success** operation to B, then B will be executed only  
3240 when both A and C have been successfully executed.
- 3241 • The multiple nodes defined by an **on\_success** construct can be executed in parallel.

- 3242
- 3243
- 3244
- 3245
- Every step that doesn't have any predecessor is considered as an initial step and can run in parallel.
  - Every step that doesn't define any successor is considered as final. When all the final nodes executions are completed then the workflow is considered as completed.


3246 **7.3.1.1.1 Example**

3247 The following example defines multiple steps and the **on\_success** relationship between them.

3248

```
topology_template:
  workflows:
    deploy:
      description: Workflow to deploy the application
      steps:
        A:
          on_success:
            - B
            - C
        B:
          on_success:
            - D
        C:
          on_success:
            - D
        D:
        E:
          on_success:
            - C
            - F
        F:
```

3249 The following schema is the visualization of the above definition in term of sequencing of the steps.



3250

### 3251 7.3.1.2 Define a sequence of activity on the same element

3252 The step definition of a TOSCA imperative workflow allows multiple activities to be defined :

3253

```

workflows:
  my_workflow:
    steps:
      create_my_node:
        target: my_node
        activities:
          - set_state: creating
          - call_operation: tosca.interfaces.node.lifecycle.Standard.create
          - set_state: created
  
```

3254 The sequence defined here defines three different activities that will be performed in a sequential way.

3255 This is just equivalent to writing multiple steps chained by an on\_success together :

3256

3257

```

workflows:
  my_workflow:
    steps:
      creating_my_node:
        target: my_node
        activities:
          - set_state: creating
        on_success: create_my_node
      create_my_node:
        target: my_node
  
```

```


activities:
  - call_operation: tosca.interfaces.node.lifecycle.Standard.create
  on_success: created_my_node
created_my_node:
  target: my_node
  activities:
    - set_state: created

```

3258

3259 In both situations the resulting workflow is a sequence of activities:

3260



3261

### 3262 7.3.2 Definition of a simple workflow

3263 Imperative workflow allow user to define custom workflows allowing them to add operations that are not  
 3264 normative, or for example, to execute some operations in parallel when TOSCA would have performed  
 3265 sequential execution.

3266

3267 As Imperative workflows are related to a topology, adding a workflow is as simple as adding a workflows  
 3268 section to your topology template and specifying the workflow and the steps that compose it.

#### 3269 7.3.2.1 Example: Adding a non-normative custom workflow

3270 This sample topology add a very simple custom workflow to trigger the mysql backup operation.

```

topology_template:
  node_templates:
    my_server:
      type: tosca.nodes.Compute
    mysql:
      type: tosca.nodes.DBMS.MySQL
      requirements:
        - host: my_server
      interfaces:
        tosca.interfaces.nodes.custom.Backup:
          operations:
            backup: backup.sh
  workflows:

```

```

backup:
  description: Performs a snapshot of the MySQL data.
  steps:
    my_step:
      target: mysql
      activities:
        - call_operation: tosca.interfaces.nodes.custom.Backup.backup


```

3271

3272 In such topology the TOSCA container will still use declarative workflow to generate the deploy and  
 3273 undeploy workflows as they are not specified and a backup workflow will be available for user to trigger.

3274 **7.3.2.2 Example: Creating two nodes hosted on the same compute in parallel**


3275 TOSCA declarative workflow generation constraint the workflow so that no operations are called in  
 3276 parallel on the same host. Looking at the following topology this means that the mysql and tomcat nodes  
 3277 will not be created in parallel but sequentially. This is fine in most of the situations as packet managers  
 3278 like apt or yum doesn't not support concurrency, however if both create operations performs a download  
 3279 of zip package from a server most of people will hope to do that in parallel in order to optimize throughput.



3280

3281 Imperative workflows can help to solve this issue. Based on the above topology we will design a workflow  
 3282 that will create tomcat and mysql in parallel but we will also ensure that tomcat is started after mysql is  
 3283 started even if no relationship is defined between the components:

3284



3285

3286

3287 To achieve such workflow, the following topology will be defined:

3288

```

topology_template:
  node_templates:
    my_server:
      type: tosca.nodes.Compute
    mysql:
      type: tosca.nodes.DBMS.MySQL
      requirements:
        - host: my_server
    tomcat:
      type: tosca.nodes.WebServer.Tomcat
      requirements:
        - host: my_server
  workflows:
    deploy:
      description: Override the TOSCA declarative workflow with the following.
      steps:
        compute_install
        target: my_server

```

```

activities:
  - delegate: deploy
on_success:
  - mysql_install
  - tomcat_install
tomcat_install:
  target: tomcat
  activities:
    - set_state: creating
    - call_operation: tosca.interfaces.node.lifecycle.Standard.create
    - set_state: created
  on_success:
    - tomcat_starting
mysql_install:
  target: mysql
  activities:
    - set_state: creating
    - call_operation: tosca.interfaces.node.lifecycle.Standard.create
    - set_state: created
    - set_state: starting
    - call_operation: tosca.interfaces.node.lifecycle.Standard.start
    - set_state: started
  on_success:
    - tomcat_starting
tomcat_starting:
  target: tomcat
  activities:
    - set_state: starting
    - call_operation: tosca.interfaces.node.lifecycle.Standard.start
    - set_state: started

```

3289

### 3290 **7.3.3 Specifying preconditions to a workflow**

3291 Pre conditions allows the TOSCA orchestrator to determine if a workflow can be executed based on the  
 3292 states and attribute values of the topology's node. Preconditions must be added to the initial workflow.

#### 3293 **7.3.3.1 Example : adding precondition to custom backup workflow**

3294 In this example we will use precondition so that we make sure that the mysql node is in the correct state  
 3295 for a backup.

```

topology_template:
  node_templates:
    my_server:
      type: toska.nodes.Compute
    mysql:
      type: toska.nodes.DBMS.MySQL
      requirements:
        - host: my_server
      interfaces:
        toska.interfaces.nodes.custom.Backup:
          operations:
            backup: backup.sh
  workflows:
    backup:
      description: Performs a snapshot of the MySQL data.
      preconditions:
        - target: my_server
          condition:
            - assert:
              - state: [{equal: available}]
        - target: mysql
          condition:
            - assert:
              - state: [{valid_values: [started, available]}]
              - my_attribute: [{equal: ready }]
      steps:
        my_step:
          target: mysql
          activities:
            - call_operation: toska.interfaces.nodes.custom.Backup.backup

```

3296 When the backup workflow will be triggered (by user or policy) the TOSCA engine will first check that  
 3297 preconditions are fulfilled. In this situation the engine will check that *my\_server* node is in *available* state  
 3298 AND that *mysql* node is in *started* OR *available* states AND that *mysql my\_attribute* value is equal to  
 3299 *ready*.

### 3300 **7.3.4 Workflow reusability**

3301 TOSCA allows the reusability of a workflow in other workflows. Such concepts can be achieved thanks to  
 3302 the inline activity.

#### 3303 **7.3.4.1 Reusing a workflow to build multiple workflows**

3304 The following example show how a workflow can inline an existing workflow and reuse it.

3305

```

topology_template:
  node_templates:
    my_server:
      type: tosca.nodes.Compute
    mysql:
      type: tosca.nodes.DBMS.MySQL
      requirements:
        - host: my_server
      interfaces:
        tosca.interfaces.nodes.custom.Backup:
          operations:
            backup: backup.sh
  workflows:
    start_mysql:
      steps:
        start_mysql:
          target: mysql
          activities :
            - set_state: starting
            - call_operation: tosca.interfaces.node.lifecycle.Standard.start
            - set_state: started
    stop_mysql:
      steps:
        stop_mysql:
          target: mysql
          activities:
            - set_state: stopping
            - call_operation: tosca.interfaces.node.lifecycle.Standard.stop
            - set_state: stopped

  backup:
    description: Performs a snapshot of the MySQL data.
    preconditions:
      - target: my_server
        condition:
          - assert:
              - state: [{equal: available}]
      - target: mysql
        condition:
          - assert:
              - state: [{valid_values: [started, available]}]

```

```


    - my_attribute: [{equal: ready }]
  steps:
    backup_step:
      activities:
        - inline: stop
        - call_operation: toska.interfaces.nodes.custom.Backup.backup
        - inline: start
  restart:
    steps:
      backup_step:
        activities:
          - inline: stop
          - inline: start

```

3306

3307 The example above defines three workflows and show how the start\_mysql and stop\_mysql workflows  
 3308 are reused in the backup and restart workflows.

3309 Inlined workflows are inlined sequentially in the existing workflow for example the backup workflow would  
 3310 look like this:



3311

### 3312 7.3.4.2 Inlining a complex workflow

3313 It is possible of course to inline more complex workflows. The following example defines an inlined  
 3314 workflows with multiple steps including concurrent steps:

3315

```

topology_template:
  workflows:
    inlined_wf:
      steps:

```


```
A:
  target: node_a
  activities:
    - call_operation: a
  on_success:
    - B
    - C
B:
  target: node_a
  activities:
    - call_operation: b
  on_success:
    - D
C:
  target: node_a
  activities:
    - call_operation: c
  on_success:
    - D
D:
  target: node_a
  activities:
    - call_operation: d
E:
  target: node_a
  activities:
    - call_operation: e
  on_success:
    - C
    - F
F:
  target: node_a
  activities:
    - call_operation: f
main_workflow:
  steps:
    G:
      target: node_a
      activities:
        - set_state: initial
        - inline: inlined_wf
```

```
- set_state: available
```

3316

3317 To describe the following workflow:

3318



3319

### 3320 7.3.5 Defining conditional logic on some part of the workflow

3321 Preconditions are used to validate if the workflow should be executed only for the initial workflow. If a  
3322 workflow that is inlined defines some preconditions these preconditions will be used at the instance level  
3323 to define if the operations should be executed or not on the defined instance.

3324

3325 This construct can be used to filter some steps on a specific instance or under some specific  
3326 circumstances or topology state.

3327

```
topology_template:
  node_templates:
    my_server:
      type: tosca.nodes.Compute
    cluster:
      type: tosca.nodes.DBMS.Cluster
    requirements:
      - host: my_server
    interfaces:
      tosca.interfaces.nodes.custom.Backup:
        operations:
          backup: backup.sh
```

```

workflows:
  backup:
    description: Performs a snapshot of the MySQL data.
    preconditions:
      - target: my_server
        condition:
          - assert:
              - state: [{equal: available}]
      - target: mysql
        condition:
          - assert:
              - state: [{valid_values: [started, available]}]
              - my_attribute: [{equal: ready }]
    steps:
      backup_step:
        target: cluster
        filter: # filter is a list of clauses. Matching between clauses is and.
          - or: # only one of sub-clauses must be true.
              - assert:
                  - foo: [{equals: true}]
              - assert:
                  - bar: [{greater_than: 2}, {less_than: 20}]
        activities:
          - call_operation: toska.interfaces.nodes.custom.Backup.backup

```

3328

### 3329 **7.3.6 Define inputs for a workflow**

3330 Inputs can be defined in a workflow and will be provided in the execution context of the workflow. If an  
 3331 operation defines a `get_input` function on one of its parameter the input will be retrieved from the workflow  
 3332 input, and if not found from the topology inputs.

3333

3334 Workflow inputs will never be configured from policy triggered workflows and SHOULD be used only for  
 3335 user triggered workflows. Of course operations can still refer to topology inputs or template properties or  
 3336 attributes even in the context of a policy triggered workflow.

#### 3337 **7.3.6.1 Example**

```

topology_template:
  node_templates:
    my_server:
      type: toska.nodes.Compute
    mysql:
      type: toska.nodes.DBMS.MySQL

```

```

requirements:
  - host: my_server
interfaces:
  tosca.interfaces.nodes.custom.Backup:
    operations:
      backup:
        implementation: backup.sh
        inputs:
          storage_url: { get_input: storage_url }
workflows:
  backup:
    description: Performs a snapshot of the MySQL data.
    preconditions:
      - target: my_server
        valid_states: [available]
      - target: mysql
        valid_states: [started, available]
    attributes:
      my_attribute: [ready]
    inputs:
      storage_url:
        type: string
    steps:
      my_step:
        target: mysql
        activities:
          - call_operation: tosca.interfaces.nodes.custom.Backup.backup

```

3338

3339 To trigger such a workflow, the TOSCA engine must allow user to provide inputs that match the given  
3340 definitions.

### 3341 **7.3.7 Handle operation failure**

3342 By default, failure of any activity of the workflow will result in the failure of the workflow and will results in  
3343 stopping the steps to be executed.

3344

3345 Exception: uninstall workflow operation failure SHOULD not prevent the other operations of the workflow  
3346 to run (a failure in an uninstall script SHOULD not prevent from releasing resources from the cloud).


3347

3348 For any workflow other than install and uninstall failures may leave the topology in an unknown state. In  
3349 such situation the TOSCA engine may not be able to orchestrate the deployment. Implementation of  
3350 **on\_failure** construct allows to execute rollback operations and reset the state of the affected entities  
3351 back to an orchestrator known state.

3352 **7.3.7.1 Example**

```
topology_template:
  node_templates:
    my_server:
      type: tosca.nodes.Compute
    mysql:
      type: tosca.nodes.DBMS.MySQL
      requirements:
        - host: my_server
      interfaces:
        tosca.interfaces.nodes.custom.Backup:
          operations:
            backup:
              implementation: backup.sh
              inputs:
                storage_url: { get_input: storage_url }
  workflows:
    backup:
      steps:
        backup_step:
          target: mysql
          activities:
            - set_state: backing_up # this state is not a TOSCA known state
            - call_operation: tosca.interfaces.nodes.custom.Backup.backup
            - set_state: available # this state is known by TOSCA orchestrator
          on_failure:
            - rollback_step
        rollback_step:
          target: mysql
          activities:
            - call_operation: tosca.interfaces.nodes.custom.Backup.backup
            - set_state: available # this state is known by TOSCA orchestrator
```

3353



3354  
3355

## 3356 7.4 Making declarative more flexible and imperative more generic

3357 TOSCA simple profile 1.1 version provides the genericity and reusability of declarative workflows that is  
3358 designed to address most of use-cases and the flexibility of imperative workflows to address more  
3359 complex or specific use-cases.

3360

3361 Each approach has some pros and cons and we are working so that the next versions of the specification  
3362 can improve the workflow usages to try to allow more flexibility in a more generic way. Two non-exclusive  
3363 leads are currently being discussed within the working group and may be included in the future versions  
3364 of the specification.

- 3365 • Improvement of the declarative workflows in order to allow people to extend the weaving logic of
- 3366 TOSCA to fit some specific need.
- 3367 • Improvement of the imperative workflows in order to allow partial imperative workflows to be
- 3368 automatically included in declarative workflows based on specific constraints on the topology
- 3369 elements.

3370 Implementation of the improvements will be done by adding some elements to the specification and will  
3371 not break compatibility with the current specification.

### 3372 7.4.1.1 Notes

- 3373 • The weaving improvement section is a Work in Progress and is not final in 1.1 version. The
- 3374 elements in this section are incomplete and may be subject to change in next specification
- 3375 version.
- 3376 • Moreover, the weaving improvements is one of the track of improvements. As describe improving
- 3377 the reusability of imperative workflow is another track (that may both co-exists in next
- 3378 specifications).

## 3379 7.4.2 Weaving improvements

3380 Making declarative better experimental option.

### 3381 7.4.2.1 Node lifecycle definition

3382 Node workflow is defined at the node type level. The node workflow definition is used to generate the  
3383 declarative workflow of a given node.

3384 The `tosca.nodes.Root` type defines workflow steps for both the install workflow (used to instantiate or  
3385 deploy a topology) and the uninstall workflow (used to destroy or undeploy a topology). The workflow is  
3386 defined as follows:  
3387

```
node_types:
  tosca.nodes.Root:
    workflows:
      install:
        steps:
          install_sequence:
            activities:
              - set_state: creating
              - call_operation: tosca.interfaces.node.lifecycle.Standard.create
              - set_state: created
              - set_state: configuring
              - call_operation:
                tosca.interfaces.node.lifecycle.Standard.configure
              - set_state: configured
              - set_state: starting
              - call_operation: tosca.interfaces.node.lifecycle.Standard.start
              - set_state: started
          uninstall:
            steps:
              uninstall_sequence:
                activities:
                  - set_state: stopping
                  - call_operation: tosca.interfaces.node.lifecycle.Standard.stop
                  - set_state: stopped
                  - set_state: deleting
                  - call_operation: tosca.interfaces.node.lifecycle.Standard.delete
                  - set_state: deleted
```

3388

### 3389 **7.4.2.2 Relationship lifecycle and weaving**

3390 While the workflow of a single node is quite simple the TOSCA weaving process is the real key element of  
3391 declarative workflows. The process of weaving consist of the ability to create complex management  
3392 workflows including dependency management in execution order between node operations, injection of  
3393 operations to process specific instruction related to the connection to other nodes based the relationships  
3394 and groups defined in a topology.

3395

3396 This section describes the relationship weaving and how the description at a template level can be  
3397 translated on an instance level.

```

relationship_types:
  tosca.relationships.ConnectsTo:
    workflow:
      install: # name of the workflow for wich the weaving has to be taken in
account
      source_weaving: # Instruct how to weave some tasks on the source workflow
(executed on SOURCE instance)
        - after: configuring # instruct that this operation should be weaved
after the target reach configuring state
          wait_target: created # add a join from a state of the target
          activity:
tosca.interfaces.relationships.Configure.pre_configure_source
        - before: configured # instruct that this operation should be weaved
before the target reach configured state
          activity:
tosca.interfaces.relationships.Configure.post_configure_source
        - before: starting
          wait_target: started # add a join from a state of the target
        - after: started
          activity: tosca.interfaces.relationships.Configure.add_target
      target_weaving: # Instruct how to weave some tasks on the target workflow
(executed on TARGET instance)
        - after: configuring # instruct that this operation should be weaved
after the target reach configuring state
          after_source: created # add a join from a state of the source
          activity:
tosca.interfaces.relationships.Configure.pre_configure_target
        - before: configured # instruct that this operation should be weaved
before the target reach configured state
          activity:
tosca.interfaces.relationships.Configure.post_configure_target
        - after: started
          activity: tosca.interfaces.relationships.Configure.add_source

```

3399

## 8 TOSCA networking

3400 Except for the examples, this section is **normative** and describes how to express and control the  
3401 application centric network semantics available in TOSCA.

### 8.1 Networking and Service Template Portability

3403 TOSCA Service Templates are application centric in the sense that they focus on describing application  
3404 components in terms of their requirements and interrelationships. In order to provide cloud portability, it is  
3405 important that a TOSCA Service Template avoid cloud specific requirements and details. However, at the  
3406 same time, TOSCA must provide the expressiveness to control the mapping of software component  
3407 connectivity to the network constructs of the hosting cloud.

3408 TOSCA Networking takes the following approach.

- 3409 1. The application component connectivity semantics and expressed in terms of Requirements and  
3410 Capabilities and the relationships between these. Service Template authors are able to express  
3411 the interconnectivity requirements of their software components in an abstract, declarative, and  
3412 thus highly portable manner.
- 3413 2. The information provided in TOSCA is complete enough for a TOSCA implementation to fulfill the  
3414 application component network requirements declaratively (i.e., it contains information such as  
3415 communication initiation and layer 4 port specifications) so that the required network semantics  
3416 can be realized on arbitrary network infrastructures.
- 3417 3. TOSCA Networking provides full control of the mapping of software component interconnectivity  
3418 to the networking constructs of the hosting cloud network independently of the Service Template,  
3419 providing the required separation between application and network semantics to preserve Service  
3420 Template portability.
- 3421 4. Service Template authors have the choice of specifying application component networking  
3422 requirements in the Service Template or completely separating the application component to  
3423 network mapping into a separate document. This allows application components with explicit  
3424 network requirements to express them while allowing users to control the complete mapping for  
3425 all software components which may not have specific requirements. Usage of these two  
3426 approaches is possible simultaneously and required to avoid having to re-write components  
3427 network semantics as arbitrary sets of components are assembled into Service Templates.
- 3428 5. Defining a set of network semantics which are expressive enough to address the most common  
3429 application connectivity requirements while avoiding dependencies on specific network  
3430 technologies and constructs. Service Template authors and cloud providers are able to express  
3431 unique/non-portable semantics by defining their own specialized network Requirements and  
3432 Capabilities.

### 8.2 Connectivity Semantics

3434 TOSCA's application centric approach includes the modeling of network connectivity semantics from an  
3435 application component connectivity perspective. The basic premise is that applications contain  
3436 components which need to communicate with other components using one or more endpoints over a  
3437 network stack such as TCP/IP, where connectivity between two components is expressed as a <source  
3438 component, source address, source port, target component, target address, target port> tuple. Note that  
3439 source and target components are added to the traditional 4 tuple to provide the application centric  
3440 information, mapping the network to the source or target component involved in the connectivity.


3441

3442 Software components are expressed as Node Types in TOSCA which can express virtually any kind of  
3443 concept in a TOSCA model. Node Types offering network based functions can model their connectivity  
3444 using a special Endpoint Capability, [tosca.capabilities.Endpoint](#), designed for this purpose. Node Types

3445 which require an Endpoint can specify this as a TOSCA requirement. A special Relationship Type,  
3446 `tosca.relationships.ConnectsTo`, is used to implicitly or explicitly relate the source Node Type's endpoint  
3447 to the required endpoint in the target node type. Since `tosca.capabilities.Endpoint` and  
3448 `tosca.relationships.ConnectsTo` are TOSCA types, they can be used in templates and extended by  
3449 subclassing in the usual ways, thus allowing the expression of additional semantics as needed.

3450 The following diagram shows how the TOSCA node, capability and relationship types enable modeling  
3451 the application layer decoupled from the network model intersecting at the Compute node using the  
3452 [Bindable](#) capability type.

3453 As you can see, the Port node type effectively acts a broker node between the Network node description



3454 and a host Compute node of an application.

## 3455 **8.3 Expressing connectivity semantics**

3456 This section describes how TOSCA supports the typical client/server and group communication  
3457 semantics found in application architectures.

### 3458 **8.3.1 Connection initiation semantics**

3459 The `tosca.relationships.ConnectsTo` expresses that requirement that a source application component  
3460 needs to be able to communicate with a target software component to consume the services of the target.  
3461 `ConnectTo` is a component interdependency semantic in the most general sense and does not try imply  
3462 how the communication between the source and target components is physically realized.

3463

3464 Application component intercommunication typically has conventions regarding which component(s)  
3465 initiate the communication. Connection initiation semantics are specified in [tosca.capabilities.Endpoint](#).  
3466 Endpoints at each end of the [tosca.relationships.ConnectsTo](#) must indicate identical connection initiation  
3467 semantics.

3468

3469 The following sections describe the normative connection initiation semantics for the  
3470 `tosca.relationships.ConnectsTo` Relationship Type.

#### 3471 **8.3.1.1 Source to Target**

3472 The Source to Target communication initiation semantic is the most common case where the source  
3473 component initiates communication with the target component in order to fulfill an instance of the  
3474 `tosca.relationships.ConnectsTo` relationship. The typical case is a “client” component connecting to a  
3475 “server” component where the client initiates a stream oriented connection to a pre-defined transport  
3476 specific port or set of ports.

3477  
3478 It is the responsibility of the TOSCA implementation to ensure the source component has a suitable  
3479 network path to the target component and that the ports specified in the respective  
3480 [tosca.capabilities.Endpoint](#) are not blocked. The TOSCA implementation may only represent state of the  
3481 `tosca.relationships.ConnectsTo` relationship as fulfilled after the actual network communication is enabled  
3482 and the source and target components are in their operational states.

3483  
3484 Note that the connection initiation semantic only impacts the fulfillment of the actual connectivity and does  
3485 not impact the node traversal order implied by the `tosca.relationships.ConnectsTo` Relationship Type.

### 3486 **8.3.1.2 Target to Source**

3487 The Target to Source communication initiation semantic is a less common case where the target  
3488 component initiates communication with the source component in order to fulfill an instance of the  
3489 `tosca.relationships.ConnectsTo` relationship. This “reverse” connection initiation direction is typically  
3490 required due to some technical requirements of the components or protocols involved, such as the  
3491 requirement that SSH must only be initiated from target component in order to fulfill the services required  
3492 by the source component.

3493  
3494 It is the responsibility of the TOSCA implementation to ensure the source component has a suitable  
3495 network path to the target component and that the ports specified in the respective  
3496 `tosca.capabilities.Endpoint` are not blocked. The TOSCA implementation may only represent state of the  
3497 `tosca.relationships.ConnectsTo` relationship as fulfilled after the actual network communication is enabled  
3498 and the source and target components are in their operational states.

3499  
3500 Note that the connection initiation semantic only impacts the fulfillment of the actual connectivity and does  
3501 not impact the node traversal order implied by the `tosca.relationships.ConnectsTo` Relationship Type.

### 3502 **8.3.1.3 Peer-to-Peer**

3503 The Peer-to-Peer communication initiation semantic allows any member of a group to initiate  
3504 communication with any other member of the same group at any time. This semantic typically appears in  
3505 clustering and distributed services where there is redundancy of components or services.

3506  
3507 It is the responsibility of the TOSCA implementation to ensure the source component has a suitable  
3508 network path between all the member component instances and that the ports specified in the respective  
3509 `tosca.capabilities.Endpoint` are not blocked, and the appropriate multicast communication, if necessary,  
3510 enabled. The TOSCA implementation may only represent state of the `tosca.relationships.ConnectsTo`  
3511 relationship as fulfilled after the actual network communication is enabled such that at least one-member  
3512 component of the group may reach any other member component of the group.

3513  
3514 Endpoints specifying the Peer-to-Peer initiation semantic need not be related with a  
3515 `tosca.relationships.ConnectsTo` relationship for the common case where the same set of component  
3516 instances must communicate with each other.

3517  
3518 Note that the connection initiation semantic only impacts the fulfillment of the actual connectivity and does  
3519 not impact the node traversal order implied by the `tosca.relationships.ConnectsTo` Relationship Type.

## 3520 **8.3.2 Specifying layer 4 ports**

3521 TOSCA Service Templates must express enough details about application component  
3522 intercommunication to enable TOSCA implementations to fulfill these communication semantics in the  
3523 network infrastructure. TOSCA currently focuses on TCP/IP as this is the most pervasive in today's cloud

3524 infrastructures. The layer 4 ports required for application component intercommunication are specified in  
3525 `tosca.capabilities.Endpoint`. The union of the port specifications of both the source and target  
3526 `tosca.capabilities.Endpoint` which are part of the `tosca.relationships.ConnectsTo` Relationship Template  
3527 are interpreted as the effective set of ports which must be allowed in the network communication.

3528

3529 The meaning of Source and Target port(s) corresponds to the direction of the respective  
3530 `tosca.relationships.ConnectsTo`.

## 3531 **8.4 Network provisioning**

### 3532 **8.4.1 Declarative network provisioning**

3533 TOSCA orchestrators are responsible for the provisioning of the network connectivity for declarative  
3534 TOSCA Service Templates (Declarative TOSCA Service Templates don't contain explicit plans). This  
3535 means that the TOSCA orchestrator must be able to infer a suitable logical connectivity model from the  
3536 Service Template and then decide how to provision the logical connectivity, referred to as "fulfillment", on  
3537 the available underlying infrastructure. In order to enable fulfillment, sufficient technical details still must  
3538 be specified, such as the required protocols, ports and QOS information. TOSCA connectivity types, such  
3539 as `tosca.capabilities.Endpoint`, provide well defined means to express these details.

### 3540 **8.4.2 Implicit network fulfillment**

3541 TOSCA Service Templates are by default network agnostic. TOSCA's application centric approach only  
3542 requires that a TOSCA Service Template contain enough information for a TOSCA orchestrator to infer  
3543 suitable network connectivity to meet the needs of the application components. Thus Service Template  
3544 designers are not required to be aware of or provide specific requirements for underlying networks. This  
3545 approach yields the most portable Service Templates, allowing them to be deployed into any  
3546 infrastructure which can provide the necessary component interconnectivity.




### 3547 **8.4.3 Controlling network fulfillment**

3548 TOSCA provides mechanisms for providing control over network fulfillment.




3549 This mechanism allows the application network designer to express in service template or network  
3550 template how the networks should be provisioned.

3551

3552 For the use cases described below let's assume we have a typical 3-tier application which is consisting of  
3553 FE (frontend), BE (backend) and DB (database) tiers. The simple application topology diagram can be  
3554 shown below:



3555



3556

3557

Figure-5: Typical 3-Tier Network

3558 **8.4.3.1 Use case: OAM Network**


3559 When deploying an application in service provider's on-premise cloud, it's very common that one or more  
 3560 of the application's services should be accessible from an ad-hoc OAM (Operations, Administration and  
 3561 Management) network which exists in the service provider backbone.

3562

3563 As an application network designer, I'd like to express in my TOSCA network template (which  
 3564 corresponds to my TOSCA service template) the network CIDR block, start ip, end ip and segmentation  
 3565 ID (e.g. VLAN id).

3566 The diagram below depicts a typical 3-tiers application with specific networking requirements for its FE  
 3567 tier server cluster:


3568



3569

### 3570 8.4.3.2 Use case: Data Traffic network

3571 The diagram below defines a set of networking requirements for the backend and DB tiers of the 3-tier  
3572 app mentioned above.




3573

### 3574 8.4.3.3 Use case: Bring my own DHCP

3575 The same 3-tier app requires for its admin traffic network to manage the IP allocation by its own DHCP  
 3576 which runs autonomously as part of application domain.

3577

3578 For this purpose, the app network designer would like to express in TOSCA that the underlying  
 3579 provisioned network will be set with DHCP\_ENABLED=false. See this illustrated in the figure below:



3580

## 3581 8.5 Network Types

### 3582 8.5.1 tosca.nodes.network.Network

3583 The TOSCA Network node represents a simple, logical network service.

<b>Shorthand Name</b>	Network
<b>Type Qualified Name</b>	tosca:Network
<b>Type URI</b>	tosca.nodes.network.Network

#### 3584 8.5.1.1 Properties

Name	Required	Type	Constraints	Description
ip_version	no	integer	valid_values: [4, 6] default: 4	The IP version of the requested network
cidr	no	string	None	The cidr block of the requested network
start_ip	no	string	None	The IP address to be used as the 1 <sup>st</sup> one in a pool of addresses derived from the cidr block full IP range

Name	Required	Type	Constraints	Description
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.
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"> <li>• If <b>network_name</b> is provided along with <b>network_id</b> they will be used to uniquely identify an existing network and not creating a new one, means all other possible properties are not allowed.</li> <li>• <b>network_name</b> should be more convenient for using. But in case that network name uniqueness is not guaranteed then one should provide a <b>network_id</b> as well.</li> </ul>
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"> <li>• Appearance of <b>network_id</b> in network template instructs the Tosca container to use an existing network instead of creating a new one.</li> <li>• <b>network_name</b> should be more convenient for using. But in case that network name uniqueness is not guaranteed then one should add a <b>network_id</b> as well.</li> <li>• <b>network_name</b> and <b>network_id</b> can be still used together to achieve both uniqueness and convenient.</li> </ul>
segmentation_id	no	string	None	A segmentation identifier in the underlying cloud infrastructure (e.g., VLAN id, GRE tunnel id). If the <b>segmentation_id</b> is specified, the <b>network_type</b> or <b>physical_network</b> 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, <b>physical_network</b> 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 <b>network_type</b> 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.

### 3585 8.5.1.2 Attributes

Name	Required	Type	Constraints	Description
segmentation_id	no	string	None	The actual <i>segmentation_id</i> that is been assigned to the network by the underlying cloud infrastructure.

### 8.5.1.3 Definition

```
tosca.nodes.network.Network:
  derived_from: tosca.nodes.Root
  properties:
    ip_version:
      type: integer
      required: false
      default: 4
      constraints:
        - valid_values: [ 4, 6 ]
    cidr:
      type: string
      required: false
    start_ip:
      type: string
      required: false
    end_ip:
      type: string
      required: false
    gateway_ip:
      type: string
      required: false
    network_name:
      type: string
      required: false
    network_id:
      type: string
      required: false
    segmentation_id:
      type: string
      required: false
    network_type:
      type: string
      required: false
    physical_network:
      type: string
      required: false
  capabilities:
    link:
      type: tosca.capabilities.network.Linkable
```

3587 **8.5.2 tosca.nodes.network.Port**

3588 The TOSCA **Port** node represents a logical entity that associates between Compute and Network normative types.

3590 The Port node type effectively represents a single virtual NIC on the Compute node instance.

<b>Shorthand Name</b>	Port
<b>Type Qualified Name</b>	tosca:Port
<b>Type URI</b>	tosca.nodes.network.Port

3591 **8.5.2.1 Properties**

Name	Required	Type	Constraints	Description
ip_address	no	string	None	Allow the user to set a fixed IP address.  Note that this address is a request to the provider which they will attempt to fulfill but may not be able to dependent on the network the port is associated with.
order	no	integer	greater_or_equal: 0 default: 0	The order of the NIC on the compute instance (e.g. eth2).  <b>Note:</b> 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 <b>is_default</b> =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.

3592 **8.5.2.2 Attributes**

Name	Required	Type	Constraints	Description
ip_address	no	string	None	The IP address would be assigned to the associated compute instance.

3593 **8.5.2.3 Definition**

```
tosca.nodes.network.Port:
  derived_from: tosca.nodes.Root
  properties:
    ip_address:
      type: string
      required: false
    order:
      type: integer
      required: true
      default: 0
      constraints:
        - greater_or_equal: 0
  is_default:
    type: boolean
    required: false
    default: false
  ip_range_start:
    type: string
    required: false
  ip_range_end:
    type: string
    required: false
  requirements:
    - link:
      capability: tosca.capabilities.network.Linkable
      relationship: tosca.relationships.network.Linksto
    - binding:
      capability: tosca.capabilities.network.Bindable
      relationship: tosca.relationships.network.BindsTo
```

3594 **8.5.3 tosca.capabilities.network.Linkable**

3595 A node type that includes the Linkable capability indicates that it can be pointed to by a  
3596 [tosca.relationships.network.Linksto](#) relationship type.

<b>Shorthand Name</b>	Linkable
<b>Type Qualified Name</b>	tosca::Linkable
<b>Type URI</b>	tosca.capabilities.network.Linkable

3597 **8.5.3.1 Properties**

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

3598 **8.5.3.2 Definition**

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

3599 **8.5.4 tosca.relationships.network.LinksTo**

3600 This relationship type represents an association relationship between Port and Network node types.

<b>Shorthand Name</b>	LinksTo
<b>Type Qualified Name</b>	tosca:LinksTo
<b>Type URI</b>	tosca.relationships.network.LinksTo

3601 **8.5.4.1 Definition**

```
tosca.relationships.network.LinksTo:
  derived_from: tosca.relationships.DependsOn
  valid_target_types: [ tosca.capabilities.network.Linkable ]
```

3602 **8.5.5 tosca.relationships.network.BindsTo**

3603 This type represents a network association relationship between Port and Compute node types.

<b>Shorthand Name</b>	network.BindsTo
<b>Type Qualified Name</b>	tosca:BindsTo
<b>Type URI</b>	tosca.relationships.network.BindsTo

3604 **8.5.5.1 Definition**

```
tosca.relationships.network.BindsTo:
  derived_from: tosca.relationships.DependsOn
  valid_target_types: [ tosca.capabilities.network.Bindable ]
```

## 3605 8.6 Network modeling approaches

### 3606 8.6.1 Option 1: Specifying a network outside the application's Service 3607 Template

3608 This approach allows someone who understands the application's networking requirements, mapping the  
3609 details of the underlying network to the appropriate node templates in the application.


3610  
3611 The motivation for this approach is providing the application network designer a fine-grained control on  
3612 how networks are provisioned and stitched to its application by the TOSCA orchestrator and underlying  
3613 cloud infrastructure while still preserving the portability of his service template. Preserving the portability  
3614 means here not doing any modification in service template but just "plug-in" the desired network  
3615 modeling. The network modeling can reside in the same service template file but the best practice should  
3616 be placing it in a separated self-contained network template file.

3617  
3618 This "pluggable" network template approach introduces a new normative node type called Port, capability  
3619 called `tosca.capabilities.network.Linkable` and relationship type called  
3620 `tosca.relationships.network.LinksTo`.

3621 The idea of the Port is to elegantly associate the desired compute nodes with the desired network nodes  
3622 while not "touching" the compute itself.

3623  
3624 The following diagram series demonstrate the plug-ability strength of this approach.


3625 Let's assume an application designer has modeled a service template as shown in Figure 1 that  
3626 describes the application topology nodes (compute, storage, software components, etc.) with their  
3627 relationships. The designer ideally wants to preserve this service template and use it in any cloud  
3628 provider environment without any change.



3629  
3630 **Figure-6: Generic Service Template**

3631 When the application designer comes to consider its application networking requirement they typically call  
3632 the network architect/designer from their company (who has the correct expertise).

3633 The network designer, after understanding the application connectivity requirements and optionally the  
3634 target cloud provider environment, is able to model the network template and plug it to the service  
3635 template as shown in Figure 2:




3636

3637

**Figure-7: Service template with network template A**

3638 When there's a new target cloud environment to run the application on, the network designer is simply  
 3639 creates a new network template B that corresponds to the new environmental conditions and provide it to  
 3640 the application designer which packs it into the application CSAR.



3641

3642

**Figure-8: Service template with network template B**

3643 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

```

3644 **8.6.2 Option 2: Specifying network requirements within the application's**  
3645 **Service Template**

3646 This approach allows the Service Template designer to map an endpoint to a logical network.

3647 The use case shown below examines a way to express in the TOSCA YAML service template a typical 3-  
3648 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 }
```

3649

---

## 3650 9 Non-normative type definitions

3651 This section defines **non-normative** types which are used only in examples and use cases in this  
3652 specification and are included only for completeness for the reader. Implementations of this specification  
3653 are not required to support these types for conformance.

### 3654 9.1 Artifact Types

3655 This section contains are non-normative Artifact Types used in use cases and examples.

#### 3656 9.1.1 `tosca.artifacts.Deployment.Image.Container.Docker`

3657 This artifact represents a Docker “image” (a TOSCA deployment artifact type) which is a binary comprised  
3658 of one or more (a union of read-only and read-write) layers created from snapshots within the underlying  
3659 Docker **Union File System**.

##### 3660 9.1.1.1 Definition

```
tosca.artifacts.Deployment.Image.Container.Docker:  
  derived_from: tosca.artifacts.Deployment.Image  
  description: Docker Container Image
```

#### 3661 9.1.2 `tosca.artifacts.Deployment.Image.VM.ISO`

3662 A Virtual Machine (VM) formatted as an ISO standard disk image.

##### 3663 9.1.2.1 Definition

```
tosca.artifacts.Deployment.Image.VM.ISO:  
  derived_from: tosca.artifacts.Deployment.Image.VM  
  description: Virtual Machine (VM) image in ISO disk format  
  mime_type: application/octet-stream  
  file_ext: [ iso ]
```

#### 3664 9.1.3 `tosca.artifacts.Deployment.Image.VM.QCOW2`

3665 A Virtual Machine (VM) formatted as a QEMU emulator version 2 standard disk image.

##### 3666 9.1.3.1 Definition

```
tosca.artifacts.Deployment.Image.VM.QCOW2:  
  derived_from: tosca.artifacts.Deployment.Image.VM  
  description: Virtual Machine (VM) image in QCOW v2 standard disk format  
  mime_type: application/octet-stream  
  file_ext: [ qcow2 ]
```

## 3667 9.2 Capability Types

3668 This section contains are non-normative Capability Types used in use cases and examples.

3669 **9.2.1 tosca.capabilities.Container.Docker**

3670 The type indicates capabilities of a Docker runtime environment (client).

<b>Shorthand Name</b>	Container.Docker
<b>Type Qualified Name</b>	tosca:Container.Docker
<b>Type URI</b>	tosca.capabilities.Container.Docker

3671 **9.2.1.1 Properties**

Name	Required	Type	Constraints	Description
version	no	<a href="#">version[]</a>	None	The Docker version capability (i.e., the versions supported by the capability).
publish_all	no	<a href="#">boolean</a>	default: false	Indicates that all ports (ranges) listed in the <i>dockerfile</i> using the <b>EXPOSE</b> keyword be published.
publish_ports	no	list of <a href="#">PortSpec</a>	None	List of ports mappings from source (Docker container) to target (host) ports to publish.
expose_ports	no	list of <a href="#">PortSpec</a>	None	List of ports mappings from source (Docker container) to expose to other Docker containers (not accessible outside host).
volumes	no	list of <a href="#">string</a>	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	<a href="#">string</a>	None	The optional identifier of an existing host resource that should be used to run this container on.
volume_id	no	<a href="#">string</a>	None	The optional identifier of an existing storage volume (resource) that should be used to create the container's mount point(s) on.

3672 **9.2.1.2 Definition**

```
tosca.capabilities.Container.Docker:
  derived_from: tosca.capabilities.Container
  properties:
    version:
      type: list
      required: false
      entry_schema: version
    publish_all:
      type: boolean
      default: false
      required: false
    publish_ports:
      type: list
      entry_schema: PortSpec
```

```

    required: false
  expose_ports:
    type: list
    entry_schema: PortSpec
    required: false
  volumes:
    type: list
    entry_schema: string
    required: false

```

### 3673 9.2.1.3 Notes

- 3674 • When the **expose\_ports** property is used, only the **source** and **source\_range** properties of  
3675 **PortSpec** would be valid for supplying port numbers or ranges, the **target** and **target\_range**  
3676 properties would be ignored.

## 3677 9.3 Node Types

3678 This section contains non-normative node types referenced in use cases and examples. All additional  
3679 Attributes, Properties, Requirements and Capabilities shown in their definitions (and are not inherited  
3680 from ancestor normative types) are also considered to be non-normative.

### 3681 9.3.1 **tosca.nodes.Database.MySQL**

#### 3682 9.3.1.1 Properties

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

#### 3683 9.3.1.2 Definition

```

tosca.nodes.Database.MySQL:
  derived_from: tosca.nodes.Database
  requirements:
    - host:
      node: tosca.nodes.DBMS.MySQL

```

### 3684 9.3.2 **tosca.nodes.DBMS.MySQL**

#### 3685 9.3.2.1 Properties

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

#### 3686 9.3.2.2 Definition

```

tosca.nodes.DBMS.MySQL:

```

```

derived_from: tosca.nodes.DBMS
properties:
  port:
    description: reflect the default MySQL server port
    default: 3306
  root_password:
    # MySQL requires a root_password for configuration
    # Override parent DBMS definition to make this property required
    required: true
capabilities:
  # Further constrain the 'host' capability to only allow MySQL databases
  host:
    valid_source_types: [ tosca.nodes.Database.MySQL ]

```

3687 **9.3.3 [tosca.nodes.WebServer.Apache](#)**

3688 **9.3.3.1 Properties**

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

3689 **9.3.3.2 Definition**

```

tosca.nodes.WebServer.Apache:
  derived_from: tosca.nodes.WebServer

```

3690 **9.3.4 [tosca.nodes.WebApplication.WordPress](#)**

3691 This section defines a non-normative Node type for the WordPress [[WordPress](#)] application.

3692 **9.3.4.1 Properties**

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

3693 **9.3.4.2 Definition**

```

tosca.nodes.WebApplication.WordPress:
  derived_from: tosca.nodes.WebApplication
  properties:
    admin_user:
      type: string
    admin_password:
      type: string
    db_host:

```

```

type: string
requirements:
  - database_endpoint:
      capability: tosca.capabilities.Endpoint.Database
      node: tosca.nodes.Database
      relationship: tosca.relationships.ConnectsTo

```

### 3694 9.3.5 tosca.nodes.WebServer.Nodejs

3695 This non-normative node type represents a Node.js [NodeJS] web application server.

#### 3696 9.3.5.1 Properties

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

#### 3697 9.3.5.2 Definition

```

tosca.nodes.WebServer.Nodejs:
  derived_from: tosca.nodes.WebServer
  properties:
    # Property to supply the desired implementation in the Github repository
    github_url:
      required: no
      type: string
      description: location of the application on the github.
      default: https://github.com/mmm/testnode.git
  interfaces:
    Standard:
      inputs:
        github_url:
          type: string

```

### 3698 9.3.6 tosca.nodes.Container.Application.Docker

#### 3699 9.3.6.1 Properties

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

#### 3700 9.3.6.2 Definition

```

tosca.nodes.Container.Application.Docker:

```

```
derived_from:
tosca.nodes.ContainerApplicationtosca.nodes.Container.Appli
cation
requirements:
- host:
  capability: toska.capabilities.Container.Docker
```

---

## 3701 10 Component Modeling Use Cases

3702 This section is **non-normative** and includes use cases that explore how to model components and their  
3703 relationships using TOSCA Simple Profile in YAML.

### 3704 10.1.1 Use Case: Exploring the HostedOn relationship using 3705 WebApplication and WebServer

3706 This use case examines the ways TOSCA YAML can be used to express a simple hosting relationship  
3707 (i.e., **HostedOn**) using the normative TOSCA **WebServer** and **WebApplication** node types defined in this  
3708 specification.

#### 3709 10.1.1.1 WebServer declares its “host” capability

3710 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 ]
```

3711 As can be seen, the **WebServer** Node Type declares its capability to “contain” (i.e., host) other nodes  
3712 using the symbolic name “**host**” and providing the Capability Type **tosca.capabilities.Container**. It  
3713 should be noted that the symbolic name of “**host**” is not a reserved word, but one assigned by the type  
3714 designer that implies at or betokens the associated capability. The **Container** capability definition also  
3715 includes a required list of valid Node Types that can be contained by this, the **WebServer**, Node Type.  
3716 This list is declared using the keyname of **valid\_source\_types** and in this case it includes only allowed  
3717 type **WebApplication**.

#### 3718 10.1.1.2 WebApplication declares its “host” requirement

3719 The **WebApplication** node type needs to be able to describe the type of capability a target node would  
3720 have to provide in order to “host” it. The normative TOSCA capability type **tosca.capabilities.Container** is  
3721 used to describe all normative TOSCA hosting (i.e., container-containee pattern) relationships. As can be  
3722 seen below, the **WebApplication** accomplishes this by declaring a requirement with the symbolic name  
3723 “**host**” with the **capability** keyname set to **tosca.capabilities.Container**.

3724 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
```

3725 **10.1.1.2.1 Notes**

- 3726
- The symbolic name “host” is not a keyword and was selected for consistent use in TOSCA normative node types to give the reader an indication of the type of requirement being referenced. A valid HostedOn relationship could still be established between WebApplicaton and WebServer in a TOSCA Service Template regardless of the symbolic name assigned to either the requirement or capability declaration.
- 3727
- 3728
- 3729
- 3730

3731 **10.1.2 Use Case: Establishing a ConnectsTo relationship to WebServer**

3732 This use case examines the ways TOSCA YAML can be used to express a simple connection  
3733 relationship (i.e., [ConnectsTo](#)) between some service derived from the [SoftwareComponent](#) Node Type,  
3734 to the normative [WebServer](#) node type defined in this specification.

3735 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
```

3736 Since the normative **WebServer** Node Type only declares one capability of type  
3737 **tosca.capabilities.Endpoint** (or **Endpoint**, its shortname alias in TOSCA) using the symbolic name  
3738 **data\_endpoint**, the **my\_web\_service** node template does not need to declare that symbolic name on its  
3739 requirement declaration. If however, the **my\_web\_server** node was based upon some other node type  
3740 that declared more than one capability of type **Endpoint**, then the **capability** keyname could be used  
3741 to supply the desired symbolic name if necessary.

3742 **10.1.2.1 Best practice**

3743 It should be noted that the best practice for designing Node Types in TOSCA should not export two  
3744 capabilities of the same type if they truly offer different functionality (i.e., different capabilities) which  
3745 should be distinguished using different Capability Type definitions.

3746 **10.1.3 Use Case: Attaching (local) BlockStorage to a Compute node**

3747 This use case examines the ways TOSCA YAML can be used to express a simple AttachesTo  
3748 relationship between a Compute node and a locally attached BlockStorage node.

3749 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
```

3750 **10.1.4 Use Case: Reusing a BlockStorage Relationship using Relationship**  
3751 **Type or Relationship Template**

3752 This builds upon the previous use case (10.1.3) to examine how a template author could attach multiple  
3753 Compute nodes (templates) to the same BlockStorage node (template), but with slightly different property  
3754 values for the AttachesTo relationship.

3755  
3756 Specifically, several notation options are shown (in this use case) that achieve the same desired result.

3757 **10.1.4.1 Simple Profile Rationale**

3758 Referencing an explicitly declared Relationship Template is a convenience of the Simple Profile that  
3759 allows template authors an entity to set, constrain or override the properties and operations as defined in  
3760 its declared (Relationship) Type much as allowed now for Node Templates. It is especially useful when a  
3761 complex Relationship Type (with many configurable properties or operations) has several logical

3762 occurrences in the same Service (Topology) Template; allowing the author to avoid configuring these  
3763 same properties and operations in multiple Node Templates.

#### 3764 **10.1.4.2 Notation Style #1: Augment AttachesTo Relationship Type directly in** 3765 **each Node Template**

3766 This notation extends the methodology used for establishing a HostedOn relationship, but allowing  
3767 template author to supply (dynamic) configuration and/or override of properties and operations.

3768  
3769 **Note:** This option will remain valid for Simple Profile regardless of other notation (copy or aliasing) options  
3770 being discussed or adopted for future versions.  
3771

```
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
```

3772

3773 **10.1.4.3 Notation Style #2: Use the ‘template’ keyword on the Node Templates to**  
3774 **specify which named Relationship Template to use**

3775 This option shows how to explicitly declare different named Relationship Templates within the Service  
3776 Template as part of a **relationship\_templates** section (which have different property values) and can  
3777 be referenced by different Compute typed Node Templates.

3778

```
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

```

3779

3780 **10.1.4.4 Notation Style #3: Using the “copy” keyname to define a similar**  
3781 **Relationship Template**

3782 How does TOSCA make it easier to create a new relationship template that is mostly the same as one  
3783 that exists without manually copying all the same information? TOSCA provides the **copy** keyname as a  
3784 convenient way to copy an existing template definition into a new template definition as a starting point or  
3785 basis for describing a new definition and avoid manual copy. The end results are cleaner TOSCA Service  
3786 Templates that allows the description of only the changes (or deltas) between similar templates.

3787 The example below shows that the Relationship Template named **storage\_attachesto\_1** provides  
3788 some overrides (conceptually a large set of overrides) on its Type which the Relationship Template  
3789 named **storage\_attachesto\_2** wants to “copy” before perhaps providing a smaller number of overrides.

```

node_templates:

my_block_storage:
  type: BlockStorage
  properties:
    size: 10

my_web_app_tier_1:
  derived_from: Compute
  requirements:
    - attachment:
      node: my_block_storage
      relationship: storage_attachesto_1

my_web_app_tier_2:
  derived_from: Compute
  requirements:
    - attachment:

```

```

    node: my_block_storage
    relationship: storage_attachesto_2

relationship_templates:
  storage_attachesto_1:
    type: MyAttachesTo
    properties:
      location: /my_data_location
    interfaces:
      some_interface_name:
        some_operation_name_1: my_script_1.sh
        some_operation_name_2: my_script_2.sh
        some_operation_name_3: my_script_3.sh

  storage_attachesto_2:
    # Copy the contents of the "storage_attachesto_1" template into this new one
    copy: storage_attachesto_1
    # Then change just the value of the location property
    properties:
      location: /some_other_data_location

relationship_types:

  MyAttachesTo:
    derived_from: AttachesTo
    interfaces:
      some_interface_name:
        some_operation:
          implementation: default_script.sh

```

3790

# 11 Application Modeling Use Cases

3791 This section is **non-normative** and includes use cases that show how to model Infrastructure-as-a-  
 3792 Service (IaaS), Platform-as-a-Service (PaaS) and complete application uses cases using TOSCA Simple  
 3793 Profile in YAML.

## 11.1 Use cases

3795 Many of the use cases listed below can be found under the following link:

3796 <https://github.com/openstack/heat-translator/tree/master/translator/tests/data>

### 11.1.1 Overview

Name	Description
<b>Compute:</b> Create a single Compute instance with a host Operating System	Introduces a TOSCA <b>Compute</b> 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.
<b>Software Component 1:</b> Automatic deployment of a Virtual Machine (VM) image artifact	Introduces the <b>SoftwareComponent</b> node type which declares software that is hosted on a <b>Compute</b> 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 <b>SoftwareComponent</b> node template and automatically deploys it to the Compute node.
<b>BlockStorage-1:</b> Attaching Block Storage to a single Compute instance	Demonstrates how to attach a TOSCA <b>BlockStorage</b> node to a <b>Compute</b> node using the normative <b>AttachesTo</b> relationship.
<b>BlockStorage-2:</b> Attaching Block Storage using a custom Relationship Type	Demonstrates how to attach a TOSCA <b>BlockStorage</b> node to a <b>Compute</b> node using a custom RelationshipType that derives from the normative <b>AttachesTo</b> relationship.
<b>BlockStorage-3:</b> Using a Relationship Template of type AttachesTo	Demonstrates how to attach a TOSCA <b>BlockStorage</b> node to a <b>Compute</b> node using a TOSCA Relationship Template that is based upon the normative <b>AttachesTo</b> Relationship Type.
<b>BlockStorage-4:</b> Single Block Storage shared by 2-Tier Application with custom AttachesTo Type and implied relationships	This use case shows 2 <b>Compute</b> instances (2 tiers) with one BlockStorage node, and also uses a custom <b>AttachesTo</b> Relationship that provides a default mount point (i.e., <b>location</b> ) which the 1 <sup>st</sup> tier uses, but the 2 <sup>nd</sup> tier provides a different mount point.
<b>BlockStorage-5:</b> Single Block Storage shared by 2-Tier Application with custom AttachesTo Type and explicit Relationship Templates	This use case is like the previous <b>BlockStorage-4</b> use case, but also creates two relationship templates (one for each tier) each of which provide a different mount point (i.e., <b>location</b> ) which overrides the default location defined in the custom Relationship Type.
<b>BlockStorage-6:</b> Multiple Block Storage attached to different Servers	This use case demonstrates how two different TOSCA <b>BlockStorage</b> nodes can be attached to two different <b>Compute</b> nodes (i.e., servers) each using the normative <b>AttachesTo</b> relationship.
<b>Object Storage 1:</b> Creating an Object Storage service	Introduces the TOSCA <b>ObjectStorage</b> node type and shows how it can be instantiated.
<b>Network-1:</b> Server bound to a new network	Introduces the TOSCA <b>Network</b> and <b>Port</b> nodes used for modeling logical networks using the <b>LinksTo</b> and <b>BindsTo</b> Relationship Types. In this use case, the template is invoked without an existing <b>network_name</b> as an input property so a new network is created using the properties declared in the Network node.

<b>Network-2:</b> Server bound to an existing network	Shows how to use a <b>network_name</b> as an input parameter to the template to allow a server to be associated with (i.e. bound to) an existing <b>Network</b> .
<b>Network-3:</b> Two servers bound to a single network	This use case shows how two servers ( <b>Compute</b> nodes) can be associated with the same <b>Network</b> node using two logical network <b>Ports</b> .
<b>Network-4:</b> Server bound to three networks	This use case shows how three logical networks ( <b>Network</b> nodes), each with its own IP address range, can be associated with the same server ( <b>Compute</b> node).
<b>WebServer-DBMS-1:</b> WordPress [WordPress] + MySQL, single instance	Shows how to host a TOSCA <b>WebServer</b> with a TOSCA <b>WebApplication</b> , <b>DBMS</b> and <b>Database Node Types</b> along with their dependent <b>HostedOn</b> and <b>ConnectsTo</b> relationships.
<b>WebServer-DBMS-2:</b> Nodejs with PayPal Sample App and MongoDB on separate instances	Instantiates a 2-tier application with <b>Nodejs</b> and its (PayPal sample) <b>WebApplication</b> on one tier which connects a MongoDB database (which stores its application data) using a <b>ConnectsTo</b> relationship.
<b>Multi-Tier-1:</b> Elasticsearch, Logstash, Kibana (ELK)	<p>Shows <b>Elasticsearch</b>, <b>Logstash</b> and <b>Kibana</b> (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 <b>Nodejs/MongoDB</b> 2-tier application as the one being monitored. The <b>collectd</b> and <b>rsyslog</b> components are added to both the <b>WebServer</b> and <b>Database</b> tiers which work to collect data for <b>Logstash</b>.</p> <p>In addition to the application tiers, a 3<sup>rd</sup> tier is introduced with <b>Logstash</b> to collect data from the application tiers. Finally a 4<sup>th</sup> tier is added to search the <b>Logstash</b> data with <b>Elasticsearch</b> and visualize it using <b>Kibana</b>.</p> <p><b>Note:</b> This use case also shows the convenience of using a single YAML macro (declared in the <b>dsl_definitions</b> section of the TOSCA Service Template) on multiple <b>Compute</b> nodes.</p>
<b>Container-1:</b> 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 <b>wordpress</b> and connects to second <b>mysql</b> 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><b>Variation 1:</b> Docker <b>Container</b> nodes (only) providing their Docker Requirements allowing platform (orchestrator) to select/provide the underlying Docker implementation (Capability).</p>

## 3798 11.1.2 Compute: Create a single Compute instance with a host Operating System

### 3800 11.1.2.1 Description

3801 This use case demonstrates how the TOSCA Simple Profile specification can be used to stand up a  
3802 single Compute instance with a guest Operating System using a normative TOSCA **Compute** node. The  
3803 TOSCA Compute node is declarative in that the service template describes both the processor and host  
3804 operating system platform characteristics (i.e., properties declared on the capability named “**os**”  
3805 sometimes called a “flavor”) that are desired by the template author. The cloud provider would attempt to  
3806 fulfill these properties (to the best of its abilities) during orchestration.


### 3807 11.1.2.2 Features

3808 This use case introduces the following TOSCA Simple Profile features:

- 3809 • A node template that uses the normative TOSCA **Compute** Node Type along with showing an  
3810 exemplary set of its properties being configured.

- 3811 • Use of the TOSCA Service Template **inputs** section to declare a configurable value the template
- 3812 user may supply at runtime. In this case, the “host” property named “num\_cpus” (of type integer)
- 3813 is declared.
  - 3814 ○ Use of a property constraint to limit the allowed integer values for the “num\_cpus”
  - 3815 property to a specific list supplied in the property declaration.
- 3816 • Use of the TOSCA Service Template **outputs** section to declare a value the template user may
- 3817 request at runtime. In this case, the property named “instance\_ip” is declared
  - 3818 ○ The “instance\_ip” output property is programmatically retrieved from the **Compute**
  - 3819 node’s “public\_address” attribute using the TOSCA Service Template-level
  - 3820 **get\_attribute** function.

### 3821 11.1.2.3 Logical Diagram



3822

### 3823 11.1.2.4 Sample YAML

```

tosca_definitions_version: tosca_simple_yaml_1_0

description: >
  TOSCA simple profile that just defines a single compute instance and selects a
  (guest) host Operating System from the Compute node’s properties. Note, this
  example does not include default values on inputs properties.

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

  node_templates:

```

```

my_server:
  type: Compute
  capabilities:
    host:
      properties:
        disk_size: 10 GB
        num_cpus: { get_input: cpus }
        mem_size: 1 GB
    os:
      properties:
        architecture: x86_64
        type: Linux
        distribution: ubuntu
        version: 12.04
  outputs:
    private_ip:
      description: The private IP address of the deployed server instance.
      value: { get_attribute: [my_server, private_address] }

```

### 3824 11.1.2.5 Notes

- 3825 • This use case uses a versioned, Linux Ubuntu distribution on the Compute node.

## 3826 11.1.3 Software Component 1: Automatic deployment of a Virtual Machine 3827 (VM) image artifact

### 3828 11.1.3.1 Description

3829 This use case demonstrates how the TOSCA SoftwareComponent node type can be used to declare  
3830 software that is packaged in a standard Virtual Machine (VM) image file format (i.e., in this case QCOW2)  
3831 and is hosted on a TOSCA Compute node (instance). In this variation, the SoftwareComponent declares  
3832 a VM image as a deployment artifact that includes its own pre-packaged operating system and software.  
3833 The TOSCA Orchestrator detects this known deployment artifact type on the SoftwareComponent node  
3834 template and automatically deploys it to the Compute node.

### 3835 11.1.3.2 Features

3836 This use case introduces the following TOSCA Simple Profile features:


- 3837 • A node template that uses the normative TOSCA **SoftwareComponent** Node Type along with  
3838 showing an exemplary set of its properties being configured.
- 3839 • Use of the TOSCA Service Template **artifacts** section to declare a Virtual Machine (VM) image  
3840 artifact type which is referenced by the **SoftwareComponent** node template.
- 3841 • The VM file format, in this case QCOW2, includes its own guest Operating System (OS) and  
3842 therefore does **not** “require” a TOSCA **OperatingSystem** capability from the TOSCA Compute  
3843 node.

3844 **11.1.3.3 Assumptions**

3845 This use case assumes the following:

- 3846 • That the TOSCA Orchestrator (working with the Cloud provider’s underlying management  
3847 services) is able to instantiate a Compute node that has a hypervisor that supports the Virtual  
3848 Machine (VM) image format, in this case QCOW2, which should be compatible with many  
3849 standard hypervisors such as XEN and KVM.
- 3850 • This is not a “bare metal” use case and assumes the existence of a hypervisor on the machine  
3851 that is allocated to “host” the Compute instance supports (e.g. has drivers, etc.) the VM image  
3852 format in this example.

3853 **11.1.3.4 Logical Diagram**



3854

3855 **11.1.3.5 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] }

```

### 3856 11.1.3.6 Notes

- 3857 • The use of the **type** keyname on the **artifact** definition (within the **my\_virtual\_machine** node
- 3858 template) to declare the ISO image deployment artifact type (i.e.,
- 3859 **tosca.artifacts.Deployment.Image.VM.ISO**) is redundant since the file extension is “.iso”
- 3860 which associated with this known, declared artifact type.
- 3861 • This use case references a filename on the **my\_vm\_image** artifact, which indicates a Linux,
- 3862 Fedora 18, x86 VM image, only as one possible example.


## 3863 11.1.4 Block Storage 1: Using the normative AttachesTo Relationship Type

### 3864 11.1.4.1 Description

3865 This use case demonstrates how to attach a TOSCA **BlockStorage** node to a **Compute** node using the

3866 normative **AttachesTo** relationship.

3867 **11.1.4.2 Logical Diagram**



3868

3869 **11.1.4.3 Sample YAML**

```

tosca_definitions_version: tosca_simple_yaml_1_0

description: >
  TOSCA simple profile with server and attached block storage using the normative
  AttachesTo Relationship Type.

topology_template:

  inputs:
    cpus:
      type: integer
      description: Number of CPUs for the server.
      constraints:
        - valid_values: [ 1, 2, 4, 8 ]
    storage_size:
      type: scalar-unit.size
      description: Size of the storage to be created.
      default: 1 GB
    storage_snapshot_id:
      type: string
      description: >
        Optional identifier for an existing snapshot to use when creating
        storage.
    storage_location:

```

```

    type: string
    description: Block storage mount point (filesystem path).

node_templates:
  my_server:
    type: Compute
    capabilities:
      host:
        properties:
          disk_size: 10 GB
          num_cpus: { get_input: cpus }
          mem_size: 1 GB
        os:
          properties:
            architecture: x86_64
            type: linux
            distribution: fedora
            version: 18.0
      requirements:
        - local_storage:
            node: my_storage
            relationship:
              type: AttachesTo
              properties:
                location: { get_input: storage_location }

    my_storage:
      type: BlockStorage
      properties:
        size: { get_input: storage_size }
        snapshot_id: { get_input: storage_snapshot_id }

    outputs:
      private_ip:
        description: The private IP address of the newly created compute instance.
        value: { get_attribute: [my_server, private_address] }
      volume_id:
        description: The volume id of the block storage instance.
        value: { get_attribute: [my_storage, volume_id] }


```

3870 **11.1.5 Block Storage 2: Using a custom AttachesTo Relationship Type**

3871 **11.1.5.1 Description**

3872 This use case demonstrates how to attach a TOSCA **BlockStorage** node to a **Compute** node using a  
3873 custom RelationshipType that derives from the normative **AttachesTo** relationship.

3874 **11.1.5.2 Logical Diagram**



3875

3876 **11.1.5.3 Sample YAML**

3877

```
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] }


```

## 3878 11.1.6 Block Storage 3: Using a Relationship Template of type AttachesTo

### 3879 11.1.6.1 Description

3880 This use case demonstrates how to attach a TOSCA **BlockStorage** node to a **Compute** node using a  
 3881 TOSCA Relationship Template that is based upon the normative **AttachesTo** Relationship Type.

### 3882 11.1.6.2 Logical Diagram



3883

### 3884 11.1.6.3 Sample YAML

3885

```

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] }

```

3886 **11.1.7 Block Storage 4: Single Block Storage shared by 2-Tier Application**  
 3887 **with custom AttachesTo and implied relationships**


3888 **11.1.7.1 Description**

3889 This use case shows 2 compute instances (2 tiers) with one BlockStorage node, and also uses a custom  
 3890 **AttachesTo** Relationship that provides a default mount point (i.e., **location**) which the 1<sup>st</sup> tier uses,  
 3891 but the 2<sup>nd</sup> tier provides a different mount point.

3892

3893 Please note that this use case assumes both Compute nodes are accessing different directories within  
 3894 the shared, block storage node to avoid collisions.

3895 **11.1.7.2 Logical Diagram**



3896

```

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] }

```

3898 **11.1.8 Block Storage 5: Single Block Storage shared by 2-Tier Application**  
3899 **with custom AttachesTo type and explicit Relationship Templates**


3900 **11.1.8.1 Description**

3901 This use case is like the Notation1 use case, but also creates two relationship templates (one for each  
3902 tier) each of which provide a different mount point (i.e., **location**) which overrides the default location  
3903 defined in the custom Relationship Type.

3904

3905 Please note that this use case assumes both Compute nodes are accessing different directories within  
3906 the shared, block storage node to avoid collisions.

3907 **11.1.8.2 Logical Diagram**



3908

3909 **11.1.8.3 Sample YAML**

```
tosca_definitions_version: tosca_simple_yaml_1_0
```

```

description: >
  TOSCA simple profile with a single Block Storage node shared by 2-Tier Application with
  custom AttachesTo Type and explicit Relationship Templates.

relationship_types:
  MyAttachesTo:
    derived_from: tosca.relationships.AttachesTo
    properties:
      location:
        type: string
        default: /default_location

topology_template:
  inputs:
    cpus:
      type: integer
      description: Number of CPUs for the server.
      constraints:
        - valid_values: [ 1, 2, 4, 8 ]
    storage_size:
      type: scalar-unit.size
      default: 1 GB
      description: Size of the storage to be created.
    storage_snapshot_id:
      type: string
      description: >
        Optional identifier for an existing snapshot to use when creating
        storage.
    storage_location:
      type: string
      description: >
        Block storage mount point (filesystem path).

  node_templates:

    my_web_app_tier_1:
      type: tosca.nodes.Compute
      capabilities:
        host:
          properties:
            disk_size: 10 GB

```

```

    num_cpus: { get_input: cpus }
    mem_size: 4096 MB
  os:
    properties:
      architecture: x86_64
      type: Linux
      distribution: Fedora
      version: 18.0
  requirements:
    - local_storage:
        node: my_storage
        relationship: storage_attachesto_1

my_web_app_tier_2:
  type: toska.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_attaches_to_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] }


```

## 3910 11.1.9 Block Storage 6: Multiple Block Storage attached to different Servers

### 3911 11.1.9.1 Description

3912 This use case demonstrates how two different TOSCA **BlockStorage** nodes can be attached to two  
 3913 different **Compute** nodes (i.e., servers) each using the normative **AttachesTo** relationship.

3914 **11.1.9.2 Logical Diagram**



3915

3916 **11.1.9.3 Sample YAML**

```

tosca_definitions_version: toska_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: toscanodes.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: toscanodes.BlockStorage

properties:

size: { get\_input: storage\_size }

snapshot\_id: { get\_input: storage\_snapshot\_id }

my\_server2:

type: toscanodes.Compute

capabilities:

host:

properties:

disk\_size: 10 GB

```

    num_cpus: { get_input: cpus }
    mem_size: 4096 MB
  os:
    properties:
      architecture: x86_64
      type: Linux
      distribution: Fedora
      version: 18.0
  requirements:
    - local_storage:
      node: my_storage2
      relationship:
        type: AttachesTo
        properties:
          location: { get_input: storage_location }
  my_storage2:
    type: tosca.nodes.BlockStorage
    properties:
      size: { get_input: storage_size }
      snapshot_id: { get_input: storage_snapshot_id }


  outputs:
    server_ip_1:
      description: The private IP address of the application's first server.
      value: { get_attribute: [my_server, private_address] }
    server_ip_2:
      description: The private IP address of the application's second server.
      value: { get_attribute: [my_server2, private_address] }
    volume_id_1:
      description: The volume id of the first block storage instance.
      value: { get_attribute: [my_storage, volume_id] }
    volume_id_2:
      description: The volume id of the second block storage instance.
      value: { get_attribute: [my_storage2, volume_id] }

```

3917 **11.1.10 Object Storage 1: Creating an Object Storage service**

3918 **11.1.10.1 Description**

3919 **11.1.10.2 Logical Diagram**



3920

3921 **11.1.10.3 Sample YAML**

```
tosca_definitions_version: tosca_simple_yaml_1_0

description: >
  Tosca template for creating an object storage service.

topology_template:
  inputs:
    objectstore_name:
      type: string


  node_templates:
    obj_store_server:
      type: tosca.nodes.ObjectStorage
      properties:
        name: { get_input: objectstore_name }
        size: 4096 MB
        maxsize: 20 GB
```

3922 **11.1.11 Network 1: Server bound to a new network**

3923 **11.1.11.1 Description**

3924 Introduces the TOSCA **Network** and **Port** nodes used for modeling logical networks using the **LinksTo** and  
3925 **BindsTo** Relationship Types. In this use case, the template is invoked without an existing `network_name`  
3926 as an input property so a new network is created using the properties declared in the Network node.

3927 **11.1.11.2 Logical Diagram**



3928

3929 **11.1.11.3 Sample YAML**

```
tosca_definitions_version: tosca_simple_yaml_1_0

description: >
  TOSCA simple profile with 1 server bound to a new network

topology_template:

  inputs:
    network_name:
      type: string
      description: Network name

  node_templates:
    my_server:
      type: tosca.nodes.Compute
      capabilities:
        host:
          properties:
            disk_size: 10 GB
            num_cpus: 1
            mem_size: 4096 MB
          os:
            properties:
              architecture: x86_64
```

```

type: Linux
distribution: CirrOS
version: 0.3.2

my_network:
  type: toska.nodes.network.Network
  properties:
    network_name: { get_input: network_name }
    ip_version: 4
    cidr: '192.168.0.0/24'
    start_ip: '192.168.0.50'
    end_ip: '192.168.0.200'
    gateway_ip: '192.168.0.1'

my_port:
  type: toska.nodes.network.Port
  requirements:
    - binding: my_server
    - link: my_network


```

3930 **11.1.12 Network 2: Server bound to an existing network**

3931 **11.1.12.1 Description**

3932 This use case shows how to use a `network_name` as an input parameter to the template to allow a server  
 3933 to be associated with an existing network.

3934 **11.1.12.2 Logical Diagram**



3935

### 11.1.12.3 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: tosca.nodes.network.Network
      properties:
        network_name: { get_input: network_name }


    my_port:
      type: tosca.nodes.network.Port
      requirements:
        - binding:
            node: my_server
        - link:
            node: my_network
```

3937 **11.1.13 Network 3: Two servers bound to a single network**

3938 **11.1.13.1 Description**

3939 This use case shows how two servers (**Compute** nodes) can be bound to the same **Network** (node) using  
3940 two logical network **Ports**.

3941 **11.1.13.2 Logical Diagram**



3942

3943 **11.1.13.3 Sample YAML**

```
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: tosca.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
```

## 3944 **11.1.14 Network 4: Server bound to three networks**

### 3945 **11.1.14.1 Description**

3946 This use case shows how three logical networks (Network), each with its own IP address range, can be  
3947 bound to with the same server (Compute node).

3948 **11.1.14.2 Logical Diagram**



3949

3950 **11.1.14.3 Sample YAML**

```
tosca_definitions_version: toska_simple_yaml_1_0

description: >
  TOSCA simple profile with 1 server bound to 3 networks

topology_template:

node_templates:
  my_server:
    type: toska.nodes.Compute
    capabilities:
      host:
        properties:
          disk_size: 10 GB
          num_cpus: 1
          mem_size: 4096 MB
      os:
        properties:
          architecture: x86_64
```

```
    type: Linux
    distribution: CirrOS
    version: 0.3.2

my_network1:
  type: tosca.nodes.network.Network
  properties:
    cidr: '192.168.1.0/24'
    network_name: net1

my_network2:
  type: tosca.nodes.network.Network
  properties:
    cidr: '192.168.2.0/24'
    network_name: net2

my_network3:
  type: tosca.nodes.network.Network
  properties:
    cidr: '192.168.3.0/24'
    network_name: net3

my_port1:
  type: tosca.nodes.network.Port
  properties:
    order: 0
  requirements:
    - binding: my_server
    - link: my_network1

my_port2:
  type: tosca.nodes.network.Port
  properties:
    order: 1
  requirements:
    - binding: my_server
    - link: my_network2

my_port3:
  type: tosca.nodes.network.Port
  properties:
```

```

order: 2
requirements:
  - binding: my_server
  - link: my_network3


```

3951 **11.1.15 WebServer-DBMS 1: WordPress + MySQL, single instance**

3952 **11.1.15.1 Description**

3953 TOSCA simple profile service showing the WordPress web application with a MySQL database hosted on  
 3954 a single server (instance).

3955 **11.1.15.2 Logical Diagram**



3956

3957 **11.1.15.3 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: toska.nodes.WebApplication.WordPress
      properties:
        context_root: { get_input: context_root }
      requirements:
        - host: webservers
        - database_endpoint: mysql_database
      interfaces:
        Standard:
          create: wordpress_install.sh
          configure:
            implementation: wordpress_configure.sh
            inputs:
              wp_db_name: { get_property: [ mysql_database, name ] }
              wp_db_user: { get_property: [ mysql_database, user ] }
              wp_db_password: { get_property: [ mysql_database, password ] }
              # In my own template, find requirement/capability, find port
property

```

```

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

mysql_database:
  type: Database
  properties:
    name: { get_input: db_name }
    user: { get_input: db_user }
    password: { get_input: db_pwd }
    port: { get_input: db_port }
  capabilities:
    database_endpoint:
      properties:
        port: { get_input: db_port }
  requirements:
    - host: mysql_dbms
  interfaces:
    Standard:
      configure: mysql_database_configure.sh

mysql_dbms:
  type: DBMS
  properties:
    root_password: { get_input: db_root_pwd }
    port: { get_input: db_port }
  requirements:
    - host: server
  interfaces:
    Standard:
      inputs:
        db_root_password: { get_property: [ mysql_dbms, root_password ] }
      create: mysql_dbms_install.sh
      start: mysql_dbms_start.sh
      configure: mysql_dbms_configure.sh

webservice:
  type: WebServer
  requirements:
    - host: server
  interfaces:
    Standard:
      create: webservice_install.sh

```

```

    start: webservers\_start.sh

server:
  type: Compute
  capabilities:
    host:
      properties:
        disk_size: 10 GB
        num_cpus: { get_input: cpus }
        mem_size: 4096 MB
    os:
      properties:
        architecture: x86_64
        type: linux
        distribution: fedora
        version: 17.0

outputs:
  website_url:
    description: URL for Wordpress wiki.
    value: { get_attribute: [server, public_address] }

```

#### 3958 [11.1.15.4 Sample scripts](#)

3959 Where the referenced implementation scripts in the example above would have the following contents

##### 3960 [11.1.15.4.1 wordpress\\_install.sh](#)

```
yum -y install wordpress
```

##### 3961 [11.1.15.4.2 wordpress\\_configure.sh](#)

```

sed -i "/Deny from All/d" /etc/httpd/conf.d/wordpress.conf
sed -i "s/Require local/Require all granted/" /etc/httpd/conf.d/wordpress.conf
sed -i s/database_name_here/name/ /etc/wordpress/wp-config.php
sed -i s/username_here/user/ /etc/wordpress/wp-config.php
sed -i s/password_here/password/ /etc/wordpress/wp-config.php
systemctl restart httpd.service

```

##### 3962 [11.1.15.4.3 mysql\\_database\\_configure.sh](#)

```

# Setup MySQL root password and create user
cat << EOF | mysql -u root --password=db_root_password

```

```
CREATE DATABASE name;
GRANT ALL PRIVILEGES ON name.* TO "user"@"localhost"
IDENTIFIED BY "password";
FLUSH PRIVILEGES;
EXIT
EOF
```

#### 3963 **11.1.15.4.4 mysql\_dbms\_install.sh**

```
yum -y install mysql mysql-server
# Use systemd to start MySQL server at system boot time
systemctl enable mysqld.service
```

#### 3964 **11.1.15.4.5 mysql\_dbms\_start.sh**

```
# Start the MySQL service (NOTE: may already be started at image boot time)
systemctl start mysqld.service
```

#### 3965 **11.1.15.4.6 mysql\_dbms\_configure**

```
# Set the MySQL server root password
mysqladmin -u root password db_root_password
```

#### 3966 **11.1.15.4.7 webserver\_install.sh**

```
yum -y install httpd
systemctl enable httpd.service
```

#### 3967 **11.1.15.4.8 webserver\_start.sh**


```
# Start the httpd service (NOTE: may already be started at image boot time)
systemctl start httpd.service
```

### 3968 **11.1.16 WebServer-DBMS 2: Nodejs with PayPal Sample App and MongoDB** 3969 **on separate instances**

#### 3970 **11.1.16.1 Description**

3971 This use case Instantiates a 2-tier application with Nodejs and its (PayPal sample) WebApplication on  
3972 one tier which connects a MongoDB database (which stores its application data) using a ConnectsTo  
3973 relationship.

3974 **11.1.16.2 Logical Diagram**



3975

3976 **11.1.16.3 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: tosca.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 ] }
              mongoddb_ip: { get_attribute: [ mongo_server, private_address ] }
          start: scriptsscripts/nodejs/start.sh

    nodejs:
      type: tosca.nodes.WebServer.Nodejs
      requirements:
        - host: app_server
      interfaces:
        Standard:
          create: scripts/nodejs/create.sh

```

```

mongo_db:
  type: toska.nodes.Database
  requirements:
    - host: mongo_dbms
  interfaces:
    Standard:
      create: create_database.sh

mongo_dbms:
  type: toska.nodes.DBMS
  requirements:
    - host: mongo_server
  properties:
    port: 27017
  interfaces:
    toska.interfaces.node.lifecycle.Standard:
      create: mongoddb/create.sh
      configure:
        implementation: mongoddb/config.sh
      inputs:
        mongoddb_ip: { get_attribute: [mongo_server, private_address] }
      start: mongoddb/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 ] }

```

3977 **11.1.16.4 Notes:**

- 3978 • Scripts referenced in this example are assumed to be placed by the TOSCA orchestrator in the  
3979 relative directory declared in TOSCA.meta of the TOSCA CSAR file.

3980 **11.1.17 Multi-Tier-1: Elasticsearch, Logstash, Kibana (ELK) use case with**  
3981 **multiple instances**

3982 **11.1.17.1 Description**


3983 TOSCA simple profile service showing the Nodejs, MongoDB, Elasticsearch, Logstash, Kibana, rsyslog  
3984 and collectd installed on a different server (instance).

3985

3986 This use case also demonstrates:

- 3987 • Use of TOSCA macros or dsl\_definitions  
3988 • Multiple **SoftwareComponents** hosted on same Compute node  
3989 • Multiple tiers communicating to each other over ConnectsTo using Configure interface.

3990 **11.1.17.2 Logical Diagram**



3991

3992 **11.1.17.3 Sample YAML**

3993 **11.1.17.3.1 Master Service Template application (Entry-Definitions)**

3994 TheThe following YAML is the primary template (i.e., the Entry-Definition) for the overall use case. The  
3995 imported YAML for the various subcomponents are not shown here for brevity.

3996

```
tosca_definitions_version: toska_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: tosca.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: tosca.nodes.WebServer.Nodejs
  requirements:
    - host: app_server
  interfaces:
    Standard:
      create: scripts/nodejs/create.sh

mongo_db:
  type: tosca.nodes.Database
  requirements:
    - host: mongo_dbms
  interfaces:
    Standard:
      create: create_database.sh

mongo_dbms:
  type: tosca.nodes.DBMS
  requirements:
    - host: mongo_server

```

```

interfaces:
  toska.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: toska.nodes.SoftwareComponent.Elasticsearch
  requirements:
    - host: elasticsearch_server
  interfaces:
    toska.interfaces.node.lifecycle.Standard:
      create: scripts/elasticsearch/create.sh
      start: scripts/elasticsearch/start.sh
logstash:
  type: toska.nodes.SoftwareComponent.Logstash
  requirements:
    - host: logstash_server
    - search_endpoint: elasticsearch
  interfaces:
    toska.interfaces.relationship.Configure:
      pre_configure_source:
        implementation: python/logstash/configure_elasticsearch.py
      input:
        elasticsearch_ip: { get_attribute: [elasticsearch_server,
ip_address] }
  interfaces:
    toska.interfaces.node.lifecycle.Standard:
      create: scripts/logstash/create.sh
      configure: scripts/logstash/config.sh
      start: scripts/logstash/start.sh

kibana:
  type: toska.nodes.SoftwareComponent.Kibana
  requirements:
    - host: kibana_server
    - search_endpoint: elasticsearch
  interfaces:

```

```

tosca.interfaces.node.lifecycle.Standard:
  create: scripts/kibana/create.sh
  configure:
    implementation: scripts/kibana/config.sh
    input:
      elasticsearch_ip: { get_attribute: [elasticsearch_server,
ip_address] }
      kibana_ip: { get_attribute: [kibana_server, ip_address] }
    start: scripts/kibana/start.sh

app_collectd:
  type: toska.nodes.SoftwareComponent.Collectd
  requirements:
    - host: app_server
    - collectd_endpoint: logstash
  interfaces:
    toska.interfaces.relationship.Configure:
      pre_configure_target:
        implementation: python/logstash/configure_collectd.py
  interfaces:
    toska.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: toska.nodes.SoftwareComponent.Rsyslog
  requirements:
    - host: app_server
    - rsyslog_endpoint: logstash
  interfaces:
    toska.interfaces.relationship.Configure:
      pre_configure_target:
        implementation: python/logstash/configure_rsyslog.py
  interfaces:
    toska.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 ] }

```

#### 3997 **11.1.17.4 Sample scripts**

3998 Where the referenced implementation scripts in the example above would have the following contents

### 3999 **11.1.18 Container-1: Containers using Docker single Compute instance** 4000 **(Containers only)**

#### 4001 **11.1.18.1 Description**


4002 This use case shows a minimal description of two Container nodes (only) providing their Docker  
4003 Requirements allowing platform (orchestrator) to select/provide the underlying Docker implementation  
4004 (Capability). Specifically, wordpress and mysql Docker images are referenced from Docker Hub.

4005

4006 This use case also demonstrates:

- 4007 • Abstract description of Requirements (i.e., Container and Docker) allowing platform to
- 4008 dynamically select the appropriate runtime Capabilities that match.
- 4009 • Use of external repository (Docker Hub) to reference image artifact.

4010 **11.1.18.2 Logical Diagram**



4011

4012 **11.1.18.3 Sample YAML**

4013 **11.1.18.3.1 Two Docker “Container” nodes (Only) with Docker Requirements**

```

tosca_definitions_version: tosca_simple_yaml_1_0

description: >
  TOSCA simple profile with wordpress, web server and mysql on the same server.

# Repositories to retrieve code artifacts from
repositories:
  docker_hub: https://registry.hub.docker.com/

topology_template:

  inputs:
    wp_host_port:
      type: integer
      description: The host port that maps to port 80 of the WordPress container.
    db_root_pwd:
      type: string
      description: Root password for MySQL.

  node_templates:
    # The MYSQL container based on official MySQL image in Docker hub

```

```

mysql_container:
  type: tosca.nodes.Container.Application.Docker
  capabilities:
    # This is a capability that would mimic the Docker -link feature
    database_link: tosca.capabilities.Docker.Link
  artifacts:
    my_image:
      file: mysql
      type: tosca.artifacts.Deployment.Image.Container.Docker
      repository: docker_hub
  interfaces:
    Standard:
      create:
        implementation: my_image
        inputs:
          db_root_password: { get_input: db_root_pwd }

# The WordPress container based on official WordPress image in Docker hub
wordpress_container:
  type: tosca.nodes.Container.Application.Docker
  requirements:
    - database_link: mysql_container
  artifacts:
    my_image:
      file: wordpress
      type: tosca.artifacts.Deployment.Image.Container.Docker
      repository: docker_hub
  interfaces:
    Standard:
      create:
        implementation: my_image
        inputs:
          host_port: { get_input: wp_host_port }

```

4014

4015

## 12 TOSCA Policies

4016 This section is **non-normative** and describes the approach TOSCA Simple Profile plans to take for policy  
4017 description with TOSCA Service Templates. In addition, it explores how existing TOSCA Policy Types  
4018 and definitions might be applied in the future to express operational policy use cases.

### 12.1 A declarative approach

4020 TOSCA Policies are a type of requirement that govern use or access to resources which can be  
4021 expressed independently from specific applications (or their resources) and whose fulfillment is not  
4022 discretely expressed in the application's topology (i.e., via TOSCA Capabilities).

4023

4024 TOSCA deems it not desirable for a declarative model to encourage external intervention for resolving  
4025 policy issues (i.e., via imperative mechanisms external to the Cloud). Instead, the Cloud provider is  
4026 deemed to be in the best position to detect when policy conditions are triggered, analyze the affected  
4027 resources and enforce the policy against the allowable actions declared within the policy itself.

#### 12.1.1 Declarative considerations

- 4029 • Natural language rules are not realistic, too much to represent in our specification; however, regular  
4030 expressions can be used that include simple operations and operands that include symbolic names  
4031 for TOSCA metamodel entities, properties and attributes.
- 4032 • Complex rules can actually be directed to an external policy engine (to check for violation) returns  
4033 true/false then policy says what to do (trigger or action).
- 4034 • Actions/Triggers could be:
  - 4035 • Autonomic/Platform corrects against user-supplied criteria
  - 4036 • External monitoring service could be utilized to monitor policy rules/conditions against metrics,  
4037 the monitoring service could coordinate corrective actions with external services (perhaps  
4038 Workflow engines that can analyze the application and interact with the TOSCA instance model).

### 12.2 Consideration of Event, Condition and Action

### 12.3 Types of policies

4041 Policies typically address two major areas of concern for customer workloads:

- 4042 • **Access Control** – assures user and service access to controlled resources are governed by  
4043 rules which determine general access permission (i.e., allow or deny) and conditional access  
4044 dependent on other considerations (e.g., organization role, time of day, geographic location, etc.).
- 4045 • **Placement** – assures affinity (or anti-affinity) of deployed applications and their resources; that is,  
4046 what is allowed to be placed where within a Cloud provider's infrastructure.
- 4047 • **Quality-of-Service** (and continuity) - assures performance of software components (perhaps  
4048 captured as quantifiable, measure components within an SLA) along with consideration for  
4049 scaling and failover.

#### 12.3.1 Access control policies

4051 Although TOSCA Policy definitions could be used to express and convey access control policies,  
4052 definitions of policies in this area are out of scope for this specification. At this time, TOSCA encourages  
4053 organizations that already have standards that express policy for access control to provide their own  
4054 guidance on how to use their standard with TOSCA.

## 4055 12.3.2 Placement policies

4056 There must be control mechanisms in place that can be part of these patterns that accept governance  
4057 policies that allow control expressions of what is allowed when placing, scaling and managing the  
4058 applications that are enforceable and verifiable in Cloud.

4059

4060 These policies need to consider the following:

- 4061 • Regulated industries need applications to control placement (deployment) of applications to  
4062 different countries or regions (i.e., different logical geographical boundaries).

### 4063 12.3.2.1 Placement for governance concerns

4064 In general, companies and individuals have security concerns along with general “loss of control” issues  
4065 when considering deploying and hosting their highly valued application and data to the Cloud. They want  
4066 to control placement perhaps to ensure their applications are only placed in datacenter they trust or  
4067 assure that their applications and data are not placed on shared resources (i.e., not co-tenanted).

4068

4069 In addition, companies that are related to highly regulated industries where compliance with government,  
4070 industry and corporate policies is paramount. In these cases, having the ability to control placement of  
4071 applications is an especially significant consideration and a prerequisite for automated orchestration.

### 4072 12.3.2.2 Placement for failover

4073 Companies realize that their day-to-day business must continue on through unforeseen disasters that  
4074 might disable instances of the applications and data at or on specific data centers, networks or servers.  
4075 They need to be able to convey placement policies for their software applications and data that mitigate  
4076 risk of disaster by assuring these cloud assets are deployed strategically in different physical locations.  
4077 Such policies need to consider placement across geographic locations as wide as countries, regions,  
4078 datacenters, as well as granular placement on a network, server or device within the same physical  
4079 datacenter. Cloud providers must be able to not only enforce these policies but provide robust and  
4080 seamless failover such that a disaster’s impact is never perceived by the end user.

## 4081 12.3.3 Quality-of-Service (QoS) policies

4082 Quality-of-Service (apart from failover placement considerations) typically assures that software  
4083 applications and data are available and performant to the end users. This is usually something that is  
4084 measurable in terms of end-user responsiveness (or response time) and often qualified in SLAs  
4085 established between the Cloud provider and customer. These QoS aspects can be taken from SLAs and  
4086 legal agreements and further encoded as performance policies associated with the actual applications  
4087 and data when they are deployed. It is assumed that Cloud provider is able to detect high utilization (or  
4088 usage load) on these applications and data that deviate from these performance policies and is able to  
4089 bring them back into compliance.

4090

## 4091 12.4 Policy relationship considerations

- 4092 • Performance policies can be related to scalability policies. Scalability policies tell the Cloud provider  
4093 exactly **how** to scale applications and data when they detect an application’s performance policy is  
4094 (or about to be) violated (or triggered).
- 4095 • Scalability policies in turn are related to placement policies which govern **where** the application and  
4096 data can be scaled to.
- 4097 • There are general “tenant” considerations that restrict what resources are available to applications  
4098 and data based upon the contract a customer has with the Cloud provider. This includes other

4099 constraints imposed by legal agreements or SLAs that are not encoded programmatically or  
4100 associated directly with actual application or data.

## 4101 **12.5 Use Cases**

4102 This section includes some initial operation policy use cases that we wish to describe using the TOSCA  
4103 metamodel. More policy work will be done in future versions of the TOSCA Simple Profile in YAML  
4104 specification.

### 4105 **12.5.1 Placement**

#### 4106 **12.5.1.1 Use Case 1: Simple placement for failover**

##### 4107 **12.5.1.1.1 Description**

4108 This use case shows a failover policy to keep at least 3 copies running in separate containers. In this  
4109 simple case, the specific containers to use (or name is not important; the Cloud provider must assure  
4110 placement separation (anti-affinity) in three physically separate containers.

##### 4111 **12.5.1.1.2 Features**

4112 This use case introduces the following policy features:

- 4113 • Simple separation on different “compute” nodes (up to discretion of provider).
- 4114 • Simple separation by region (a logical container type) using an allowed list of region names  
4115 relative to the provider.
  - 4116 ○ Also, shows that set of allowed “regions” (containers) can be greater than the number of  
4117 containers requested.

##### 4118 **12.5.1.1.3 Logical Diagram**

4119 Sample YAML: Compute separation

```
failover_policy_1:  
  type: toska.policy.placement.Antilocate  
  description: My placement policy for Compute node separation  
  properties:  
    # 3 diff target containers  
    container type: Compute  
    container_number: 3
```

##### 4120 **12.5.1.1.4 Notes**

- 4121 • There may be availability (constraints) considerations especially if these policies are applied to  
4122 “clusters”.
- 4123 • There may be future considerations for controlling max # of instances per container.

### 4124 **12.5.1.2 Use Case 2: Controlled placement by region**

#### 4125 **12.5.1.2.1 Description**

4126 This use case demonstrates the use of named “containers” which could represent the following:

- 4127 • Datacenter regions

- 4128 • Geographic regions (e.g., cities, municipalities, states, countries, etc.)
- 4129 • Commercial regions (e.g., North America, Eastern Europe, Asia Pacific, etc.)

### 4130 **12.5.1.2.2 Features**

4131 This use case introduces the following policy features:

- 4132 • Separation of resources (i.e., TOSCA nodes) by logical regions, or zones.

### 4133 **12.5.1.2.3 Sample YAML: Region separation amongst named set of regions**

```
failover_policy_2:
  type: toska.policy.placement
  description: My failover policy with allowed target regions (logical
containers)
  properties:
    container type: region
    container_number: 3
    # If "containers" keyname is provided, they represent the allowed set
    # of target containers to use for placement for .
    containers: [ region1, region2, region3, region4 ]
```

### 4134 **12.5.1.3 Use Case 3: Co-locate based upon Compute affinity**

#### 4135 **12.5.1.3.1 Description**

4136 Nodes that need to be co-located to achieve optimal performance based upon access to similar  
4137 Infrastructure (IaaS) resource types (i.e., Compute, Network and/or Storage).

4138

4139 This use case demonstrates the co-location based upon Compute resource affinity; however, the same  
4140 approach could be taken for Network as or Storage affinity as well. :

#### 4141 **12.5.1.3.2 Features**

4142 This use case introduces the following policy features:

- 4143 • Node placement based upon Compute resource affinity.

#### 4144 **12.5.1.4 Notes**

- 4145 • The concept of placement based upon IaaS resource utilization is not future-thinking, as Cloud  
4146 should guarantee equivalent performance of application performance regardless of placement.  
4147 That is, all network access between application nodes and underlying Compute or Storage should  
4148 have equivalent performance (e.g., network bandwidth, network or storage access time, CPU  
4149 speed, etc.).

4150 **12.5.1.4.1 Sample YAML: Region separation amongst named set of regions**

```
keep_together_policy:
  type: toska.policy.placement.Colocate
  description: Keep associated nodes (groups of nodes) based upon Compute
  properties:
    affinity: Compute
```

4151 **12.5.2 Scaling**

4152 **12.5.2.1 Use Case 1: Simple node autoscale**

4153 **12.5.2.1.1 Description**

4154 Start with X nodes and scale up to Y nodes, capability to do this from a dashboard for example.

4155 **12.5.2.1.2 Features**

4156 This use case introduces the following policy features:

- 4157
  - Basic autoscaling policy

4158 **12.5.2.1.3 Sample YAML**

```
my_scaling_policy_1:
  type: toska.policy.scaling
  description: Simple node autoscaling
  properties:
    min_instances: <integer>
    max_instances: <integer>
    default_instances: <integer>
    increment: <integer>
```

4159 **12.5.2.1.4 Notes**

- 4160
  - Assume horizontal scaling for this use case
- 4161
  - Horizontal scaling, implies “stack-level” control using Compute nodes to define a “stack” (i.e., The Compute node’s entire HostedOn relationship dependency graph is considered part of its “stack”)
- 4162
- 4163
- 4164
  - Assume Compute node has a SoftwareComponent that represents a VM application.
- 4165
  - Availability Zones (and Regions if not same) need to be considered in further
- 4166 use cases.
- 4167
  - If metrics are introduced, there is a control-loop (that monitors). Autoscaling is a special concept
- 4168 that includes these considerations.
- 4169
  - Mixed placement and scaling use cases need to be considered:
- 4170
  - *Example:* Compute1 and Compute2 are 2 node templates. Compute1 has 10 instances, 5
- 4171 in one region 5 in other region.

4172

## 13 Conformance

4173

### 13.1 Conformance Targets

4174

The implementations subject to conformance are those introduced in Section 11.3 “Implementations”.

4175

They are listed here for convenience:

4176

- TOSCA YAML service template

4177

- TOSCA processor

4178

- TOSCA orchestrator (also called orchestration engine)

4179

- TOSCA generator

4180

- TOSCA archive

4181

### 13.2 Conformance Clause 1: TOSCA YAML service template

4182

A document conforms to this specification as TOSCA YAML service template if it satisfies all the

4183

statements below:

4184

- (a) It is valid according to the grammar, rules and requirements defined in section 3 “TOSCA Simple Profile definitions in YAML”.

4185

4186

- (b) When using functions defined in section 4 “TOSCA functions”, it is valid according to the grammar specified for these functions.

4187

4188

- (c) When using or referring to data types, artifact types, capability types, interface types, node types, relationship types, group types, policy types defined in section 5 “TOSCA normative type definitions”, it is valid according to the definitions given in section 5.

4189

4190

4191

### 13.3 Conformance Clause 2: TOSCA processor

4192

A processor or program conforms to this specification as TOSCA processor if it satisfies all the

4193

statements below:

4194

- (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.

4195

4196

- (b) It implements the requirements and semantics associated with the definitions and grammar in section 3 “TOSCA Simple Profile definitions in YAML”, including those listed in the “additional requirements” subsections.

4197

4198

4199

- (c) It resolves the imports, either explicit or implicit, as described in section 3 “TOSCA Simple Profile definitions in YAML”.

4200

4201

- (d) It generates errors as required in error cases described in sections 3.1 (TOSCA Namespace URI and alias), 3.2 (Parameter and property type) and 3.6 (Type-specific definitions).

4202

4203

4204

- (e) It normalizes string values as described in section 5.4.9.3 (Additional Requirements)

4205

4206

### 13.4 Conformance Clause 3: TOSCA orchestrator

4207

A processor or program conforms to this specification as TOSCA orchestrator if it satisfies all the

4208

statements below:

4209

- (a) It is conforming as a TOSCA Processor as defined in conformance clause 2: TOSCA Processor.

4210

- (b) It can process all types of artifact described in section 5.3 “Artifact types” according to the rules and grammars in this section.

4211

4212

- (c) It can process TOSCA archives as intended in section 6 “TOSCA Cloud Service Archive (CSAR) format” and other related normative sections.

4213

- 4214 (d) It can understand and process the functions defined in section 4 “TOSCA functions” according to
- 4215 their rules and semantics.
- 4216 (e) It can understand and process the normative type definitions according to their semantics and
- 4217 requirements as described in section 5 “TOSCA normative type definitions”.
- 4218 (f) It can understand and process the networking types and semantics defined in section 7 “TOSCA
- 4219 Networking”.
- 4220 (g) It generates errors as required in error cases described in sections 2.10 (Using node template
- 4221 substitution for chaining subsystems), 5.4 (Capabilities Types) and 5.7 (Interface Types).

### 4222 **13.5 Conformance Clause 4: TOSCA generator**

4223 A processor or program conforms to this specification as TOSCA generator if it satisfies at least one of  
4224 the statements below:

- 4225 (a) When requested to generate a TOSCA service template, it always produces a conforming
- 4226 TOSCA service template, as defined in Clause 1: TOSCA YAML service template,
- 4227 (b) When requested to generate a TOSCA archive, it always produces a conforming TOSCA archive,
- 4228 as defined in Clause 5: TOSCA archive.

### 4229 **13.6 Conformance Clause 5: TOSCA archive**

4230 A package artifact conforms to this specification as TOSCA archive if it satisfies all the statements below:

- 4231 (a) It is valid according to the structure and rules defined in section 6 “TOSCA Cloud Service Archive
- 4232 (CSAR) format”.

4233

## Appendix A. Known Extensions to TOSCA v1.0

4234 The following items will need to be reflected in the TOSCA (XML) specification to allow for isomorphic  
4235 mapping between the XML and YAML service templates.

### 4236 A.1 Model Changes

- 4237 • The “TOSCA Simple ‘Hello World’” example introduces this concept in Section 2. Specifically, a VM  
4238 image assumed to be accessible by the cloud provider.
- 4239 • Introduce template Input and Output parameters
- 4240 • The “Template with input and output parameter” example introduces concept in Section 2.1.1.  
4241 • “Inputs” could be mapped to BoundaryDefinitions in TOSCA v1.0. Maybe needs some usability  
4242 enhancement and better description.
- 4243 • “outputs” are a new feature.
- 4244 • Grouping of Node Templates  
4245 • This was part of original TOSCA proposal, but removed early on from v1.0 This allows grouping  
4246 of node templates that have some type of logically managed together as a group (perhaps to  
4247 apply a scaling or placement policy).
- 4248 • Lifecycle Operation definition independent/separate from Node Types or Relationship types (allows  
4249 reuse). For now, we added definitions for “node.lifecycle” and “relationship.lifecycle”.
- 4250 • Override of Interfaces (operations) in the Node Template.
- 4251 • Service Template Naming/Versioning  
4252 • Should include TOSCA spec. (or profile) version number (as part of namespace)
- 4253 • Allow the referencing artifacts using a URL (e.g., as a property value).
- 4254 • Repository definitions in Service Template.
- 4255 • Substitution mappings for Topology template.
- 4256 • Addition of Group Type, Policy Type, Group def., Policy def. along with normative TOSCA base types  
4257 for policies and groups.

### 4258 A.2 Normative Types

- 4259 • Constraints  
4260 • constraint clauses, regex
- 4261 • Types / Property / Parameters  
4262 • list, map, range, scalar-unit types  
4263 • Includes YAML intrinsic types  
4264 • NetworkInfo, PortInfo, PortDef, PortSpec, Credential  
4265 • TOSCA Version based on Maven
- 4266 • Node  
4267 • Root, Compute, ObjectStorage, BlockStorage, Network, Port, SoftwareComponent,  
4268 WebServer, WebApplicaton, DBMS, Database, Container, and others
- 4269 • Relationship  
4270 • Root, DependsOn, HostedOn, ConnectsTo, AttachesTo, RoutesTo, BindsTo, LinksTo and  
4271 others
- 4272 • Artifact  
4273 • Deployment: Image Types (e.g., VM, Container), ZIP, TAR, etc.  
4274 • Implementation: File, Bash, Python, etc.
- 4275 • Requirements

- 4276
  - None
- 4277
  - Capabilities
- 4278
  - Container, Endpoint, Attachment, Scalable, ...
- 4279
  - Lifecycle
- 4280
  - Standard (for Node Types)
- 4281
  - Configure (for Relationship Types)
- 4282
  - Functions
- 4283
  - get\_input, get\_attribute, get\_property, get\_nodes\_of\_type, get\_operation\_output and others
- 4284
  - concat, token
- 4285
  - get\_artifact
- 4286
  - Groups
- 4287
  - Root
- 4288
  - Policies
- 4289
  - Root, Placement, Scaling, Update, Performance
- 4290
  - Workflow
- 4291

4292

## Appendix B. Acknowledgments

4293 The following individuals have participated in the creation of this specification and are gratefully  
4294 acknowledged:

4295 **Contributors:**

4296 Avi Vachnis ([avi.vachnis@alcatel-lucent.com](mailto:avi.vachnis@alcatel-lucent.com)), Alcatel-Lucent

4297 Chris Lauwers ([lauwers@ubicity.com](mailto:lauwers@ubicity.com))

4298 Derek Palma ([dpalma@vnomi.com](mailto:dpalma@vnomi.com)), Vnomic

4299 Frank Leymann ([Frank.Leymann@informatik.uni-stuttgart.de](mailto:Frank.Leymann@informatik.uni-stuttgart.de)), Univ. of Stuttgart

4300 Gerd Breiter ([gbreiter@de.ibm.com](mailto:gbreiter@de.ibm.com)), IBM

4301 Hemal Surti ([hsurti@cisco.com](mailto:hsurti@cisco.com)), Cisco

4302 Ifat Afek ([ifat.afek@alcatel-lucent.com](mailto:ifat.afek@alcatel-lucent.com)), Alcatel-Lucent

4303 Idan Moyal, ([idan@gigaspaces.com](mailto:idan@gigaspaces.com)), Gigaspaces

4304 Jacques Durand ([jdurand@us.fujitsu.com](mailto:jdurand@us.fujitsu.com)), Fujitsu

4305 Jin Qin, ([chin.qinjin@huawei.com](mailto:chin.qinjin@huawei.com)), Huawei

4306 Jeremy Hess, ([jeremy@gigaspaces.com](mailto:jeremy@gigaspaces.com)), Gigaspaces

4307 John Crandall, ([mailto:jcrandal@brocade.com](mailto:mailto:jcrandal@brocade.com)), Brocade

4308 Juergen Meynert ([juergen.meynert@ts.fujitsu.com](mailto:juergen.meynert@ts.fujitsu.com)), Fujitsu

4309 Kapil Thangavelu ([kapil.thangavelu@canonical.com](mailto:kapil.thangavelu@canonical.com)), Canonical

4310 Karsten Beins ([karsten.beins@ts.fujitsu.com](mailto:karsten.beins@ts.fujitsu.com)), Fujitsu

4311 Kevin Wilson ([kevin.l.wilson@hp.com](mailto:kevin.l.wilson@hp.com)), HP

4312 Krishna Raman ([kraman@redhat.com](mailto:kraman@redhat.com)), Red Hat

4313 Luc Boutier ([luc.boutier@fastconnect.fr](mailto:luc.boutier@fastconnect.fr)), FastConnect

4314 Luca Gioppo, ([luca.gioppo@csi.it](mailto:luca.gioppo@csi.it)), CSI-Piemonte

4315 Matt Rutkowski ([mrutkows@us.ibm.com](mailto:mrutkows@us.ibm.com)), IBM

4316 Moshe Elisha ([moshe.elisha@alcatel-lucent.com](mailto:moshe.elisha@alcatel-lucent.com)), Alcatel-Lucent

4317 Nate Finch ([nate.finch@canonical.com](mailto:nate.finch@canonical.com)), Canonical

4318 Nikunj Nemani ([nnemani@vmware.com](mailto:nnemani@vmware.com)), VMware

4319 Richard Probst ([richard.probst@sap.com](mailto:richard.probst@sap.com)), SAP AG

4320 Sahdev Zala ([spzala@us.ibm.com](mailto:spzala@us.ibm.com)), IBM

4321 Shitao li ([lishitao@huawei.com](mailto:lishitao@huawei.com)), Huawei

4322 Simeon Monov ([sdmonov@us.ibm.com](mailto:sdmonov@us.ibm.com)), IBM

4323 Sivan Barzily, ([sivan@gigaspaces.com](mailto:sivan@gigaspaces.com)), Gigaspaces

4324 Sridhar Ramaswamy ([sramasw@brocade.com](mailto:sramasw@brocade.com)), Brocade

4325 Stephane Maes ([stephane.maes@hp.com](mailto:stephane.maes@hp.com)), HP

4326 Thomas Spatzier ([thomas.spatzier@de.ibm.com](mailto:thomas.spatzier@de.ibm.com)), IBM

4327 Ton Ngo ([ton@us.ibm.com](mailto:ton@us.ibm.com)), IBM

4328 Travis Tripp ([travis.tripp@hp.com](mailto:travis.tripp@hp.com)), HP

4329 Vahid Hashemian ([vahidhashemian@us.ibm.com](mailto:vahidhashemian@us.ibm.com)), IBM

4330 Wayne Witzel ([wayne.witzel@canonical.com](mailto:wayne.witzel@canonical.com)), Canonical

4331 Yaron Parasol ([yaronpa@gigaspaces.com](mailto:yaronpa@gigaspaces.com)), Gigaspaces

## Appendix C. Revision History

Revision	Date	Editor	Changes Made
WD01, Rev01	2016-01-06	Matt Rutkowski, IBM	<ul style="list-style-type: none"> <li>• Initial WD01, Revision 01 baseline for TOSCA Simple Profile in YAML v1.1</li> <li>• Cha. 10 Removed URL column for use cases in favor of a single link to Git directory where they can be found.</li> <li>• Metadata added to top-level entities</li> <li>• Policy grammar/schema fully defined.</li> <li>• Ch5. Defined TOSCA Entity Root type which is now the parent type for all TOSCA top-level types (i.e., Artifact, Capability, Relationship, Node, Group, Policy, etc.). Updated all top-level definitions to reflect in "derived_from" keyname.</li> <li>• Added TimeInterval Data Type</li> <li>• 3.5.16.1: Added keyname "schedule".</li> </ul>
WD01, Rev02	2016-01-25	Matt Rutkowski, IBM	<ul style="list-style-type: none"> <li>• 5: Removed tosca.Root type from chapter 5 until ad-hoc can agree on use cases likely to come from the TOSCA instance model WG.</li> <li>• Cleaned up TOSCA Entity Root Reorganization.</li> </ul>
WD01, Rev03	2016-03-22	Matt Rutkowski, IBM	<ul style="list-style-type: none"> <li>• 3.5.7, 3.9.3: Fixed "import" grammar (section 3.5.7) and reference to it in repository example (section 3.9.3.9.3)</li> <li>• 3.6.11.2 – Group Type – clarified group types could have members that were other groups types.</li> <li>• 5.2.5: Fixed NetworkInfo (section 5.2.5) example which was missing the 'properties' keyword.</li> <li>• 5.2.6: Clarified PortDef examples (section 5.2.6)</li> <li>• 5.2.7: Fixed PortSpec (section 5.2.7) definition to assure that target, target_range, source and source_range properties were not 'required' in schema.</li> <li>•</li> <li>• Fixed the following issues raised by TC Admin.: <ul style="list-style-type: none"> <li>• <i>Margins should be 1" top, 1" right and left, 0.5" bottom. [this will center the "new" footer, which is currently offset].</i></li> <li>• <i>The footer uses different font size (Arial 10 instead of Arial 8) and wording ("Standards Track Draft" instead of "Standards Track Work Product").</i></li> <li>• <i>Set the following three styles to use Arial 10:</i> <ul style="list-style-type: none"> <li>• "Normal around table"</li> <li>• "List Paragraph"</li> <li>• "List Bullet 3"</li> </ul> </li> <li>• <i>Around section 2.10.1, we corrected some text in the wrong font by re-applying the "normal" style.</i></li> <li>• <i>In Section 1.8 Glossary that "Node Template" definition starts off with "Relationship Template" Is that correct? Also, the paragraph formatting of the definitions seems to use weird indenting.</i></li> <li>• <i>In section 5.7.4.4, the diagram overlays the footer. We fixed this on our side by setting the preceding paragraph attribute to "keep with next".</i></li> <li>• <i>In section 2.10.2, second paragraph after Figure 1, there is a reference to "Section 0". The link jumps to 2.9.2. Is this correct?</i></li> <li>• <i>The table of examples is labelled Table of Figures. Also, the paragraph styles of these two titles should be changed from "Body text" to "Level 1", so they will show up in the TOC.</i></li> </ul> </li> <li>• 3.6.5 Interface Type – missing "derived_from" in keyname table, grammar and example.</li> </ul>
WD01, Rev04	2016-03-23	Matt Rutkowski, IBM	<ul style="list-style-type: none"> <li>• 5.2: Added section discussing TOSCA normative type names, their treatment and requirements to respect case (i.e., be case sensitive) when processing.</li> </ul>

			<ul style="list-style-type: none"> <li>• 3.6: All data types that are entity types now have their keyname tables reference the common keynames listed in section 3.6.1 TOSCA Entity schema.</li> <li>• 3.6.11: Added attributes, requirements and capabilities keynames to Group Type making it more like a Node Type (no artifacts, still logical aggregator of a set of nodes).</li> <li>• 5.9.11: Added the “network” (i.e, Endpoint) and “storage” (i.e., Storage) capabilities to the Container.Application node type.</li> </ul>
WDO1, Rev05	2016-04-20	Matt Rutkowski, IBM	<ul style="list-style-type: none"> <li>• 3.1: Bumped version number to 1.1</li> <li>• 5.3.2: typo. ‘userh’ -&gt; ‘user’ in keyname table</li> <li>• 3.6.4.4 Artifact Type - Added a note regarding “mime types” to reference official list of Apache mime types to give reader a suitable reference for expected values.</li> <li>•</li> </ul>
WDO1, Rev06	2016-17-05	Luc Boutier, FastConnect	<ul style="list-style-type: none"> <li>• 3.5.14.2.2 replaced Node Type by Node Template.</li> <li>• 3.5.17: Add workflow activity definition</li> <li>• 3.5.18: Add workflow precondition definition</li> <li>• 3.5.19: Add workflow step definition</li> <li>• 3.7.7.: Add Imperative workflow definition</li> <li>• 3.8: Add the workflows keyname to the topology template definition</li> <li>• 6.3: Added a simplified way to declare a CSAR without the meta file.</li> <li>• 7: Added a TOSCA workflows section.</li> </ul>
WDO1, Rev07	2016-19-05	Luc Boutier, FastConnect	<ul style="list-style-type: none"> <li>• 3.5.18: Add assertion definition</li> <li>• 3.5.19: Add condition clause definition</li> <li>• 3.5.20: Leverage condition clause in precondition definition</li> <li>• 3.5.21: Leverage condition clause as filter in step definition</li> <li>• 7.2: Add documentation and example on TOSCA normative weaving</li> <li>• 7.3: Fix examples in imperative workflows definition</li> </ul>
WDO1, Rev08	2016-31-05	Luc Boutier, FastConnect	<ul style="list-style-type: none"> <li>• 7.2: Specifies current expected declarative workflows and limitations.</li> </ul>
WDO1, Rev09	2016-31-05	Luc Boutier, FastConnect; Matt Rutkowski, IBM	<ul style="list-style-type: none"> <li>• 1.8: Add description for abstract nodes and no-op nodes to the glossary</li> <li>• Fixed typos, spelling/grammar and fixed numerous broken hyperlinks.</li> </ul>

4333