



TOSCA Simple Profile in YAML Version 1.0

Committee Specification Draft 03

14 April 2015

Specification URIs

This version:

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

Previous version:

<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/csd02/TOSCA-Simple-Profile-YAML-v1.0-csd02.pdf> (Authoritative)
<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/csd02/TOSCA-Simple-Profile-YAML-v1.0-csd02.html>
<http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/csd02/TOSCA-Simple-Profile-YAML-v1.0-csd02.doc>

Latest version:

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

Technical Committee:

[OASIS Topology and Orchestration Specification for Cloud Applications \(TOSCA\) TC](#)

Chairs:

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

Editors:

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

Related work:

This specification is related to:

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

Declared XML namespaces:

- <http://docs.oasis-open.org/tosca/ns/simple/yaml/1.0>

Abstract:

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

Status:

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

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

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

Citation format:

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

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

TOSCA Simple Profile in YAML Version 1.0. Edited by Derek Palma, Matt Rutkowski, and Thomas Spatzier. 14 April 2015. OASIS Committee Specification Draft 03. <http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/csd03/TOSCA-Simple-Profile-YAML-v1.0-csd03.html>. Latest version: <http://docs.oasis-open.org/tosca/TOSCA-Simple-Profile-YAML/v1.0/TOSCA-Simple-Profile-YAML-v1.0.html>.

Notices

Copyright © OASIS Open 2015. All Rights Reserved.

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

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

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

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

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

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

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

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

Table of Contents

1	Objective.....	7
2	Summary of key TOSCA concepts.....	8
3	A “hello world” template for TOSCA Simple Profile in YAML.....	9
3.1	Requesting input parameters and providing output.....	10
4	TOSCA template for a simple software installation.....	12
5	Overriding behavior of predefined node types.....	14
6	TOSCA template for database content deployment.....	16
7	TOSCA template for a two-tier application.....	19
8	Using a custom script to establish a relationship in a template.....	22
9	Using custom relationship types in a TOSCA template.....	24
9.1	Definition of a custom relationship type.....	25
10	Defining generic dependencies between nodes in a template.....	26
11	Describing abstract requirements for nodes and capabilities in a TOSCA template.....	27
11.1	Using a node_filter to define hosting infrastructure requirements for a software.....	27
11.2	Using an abstract node template to define infrastructure requirements for software.....	29
11.3	Using a node_filter to define requirements on a database for an application.....	29
12	Using node template substitution for model composition.....	31
12.1	Understanding node template instantiation through a TOSCA Orchestrator.....	31
12.2	Definition of the top-level service template.....	31
12.3	Definition of the database stack in a service template.....	33
13	Using node template substitution for chaining subsystems.....	36
13.1	Defining the overall subsystem chain.....	36
13.2	Defining a subsystem (node) type.....	38
13.3	Defining the details of a subsystem.....	39
14	Grouping node templates.....	43
15	Using YAML Macros to simplify templates.....	46
16	Passing information as inputs to Nodes and Relationships.....	47
16.1	Example: declaring input variables for all operations on a single interface.....	47
16.2	Example: declaring input variables for a single operation.....	47
16.3	Example: setting output variables to an attribute.....	48
16.4	Example: passing output variables between operations.....	48
17	Topology Template Model versus Instance Model.....	49
18	Using attributes implicitly reflected from properties.....	50
Appendix A. TOSCA Simple Profile definitions in YAML.....		52
A.1	TOSCA namespace and alias.....	52
A.2	Parameter and property types.....	52
A.3	Normative values.....	61
A.4	TOSCA Metamodel.....	62
A.5	Reusable modeling definitions.....	63
A.6	Type-specific definitions.....	79
A.7	Template-specific definitions.....	92
A.8	Topology Template definition.....	101
A.9	Service Template definition.....	106

Appendix B. Functions	115
B.1 Reserved Function Keywords	115
B.2 Environment Variable Conventions	115
B.3 Intrinsic functions	118
B.4 Property functions	119
B.5 Attribute functions	121
B.6 Operation functions	123
B.7 Navigation functions	123
B.8 Artifact functions	124
B.9 Context-based Entity name (global)	126
Appendix C. TOSCA normative type definitions	127
C.1 Assumptions	127
C.2 Data Types	127
C.3 Capabilities Types	134
C.4 Requirement Types	142
C.5 Relationship Types	142
C.6 Interface Types	146
C.7 Node Types	151
C.8 Artifact Types	162
Appendix D. Non-normative type definitions	166
D.1 Artifact Types	166
D.2 Capability Types	166
D.3 Node Types	168
Appendix E. TOSCA Cloud Service Archive (CSAR) Format	171
E.1 Overall Structure of a CSAR	171
E.2 TOSCA Meta File	171
Appendix F. Networking	172
F.1 Networking and Service Template Portability	172
F.2 Connectivity Semantics	172
F.3 Expressing connectivity semantics	173
F.4 Network provisioning	175
F.5 Network Types	179
F.6 Network modeling approaches	184
Appendix G. Component Modeling Use Cases	190
Appendix H. Complete Application Modeling Use Cases	197
H.1 Use cases	197
Appendix I. Policies (Placeholder)	242
I.1 Types of policies	242
Appendix J. References	244
J.1 Known Extensions to TOSCA v1.0	244
J.2 Terminology	245
J.3 Normative References	245
J.4 Non-Normative References	245
J.5 Glossary	245
Appendix K. Issues List	247

Appendix L. Acknowledgments	251
Appendix M. Revision History	252

Table of Examples

Example 1 - TOSCA Simple "Hello World"	9
Example 2 - Template with input and output parameter sections.....	10
Example 3 - Simple (MySQL) software installation on a TOSCA Compute node	12
Example 4 - Node Template overriding its Node Type's "configure" interface	14
Example 5 - Template for deploying database content on-top of MySQL DBMS middleware	16
Example 6 - Basic two-tier application (web application and database server tiers).....	19
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.....	26
Example 11 - Grouping Node Templates with same scaling policy	43

Table of Figures

Figure 1: Using template substitution to implement a database tier.....	32
Figure 2: Substitution mappings	34
Figure 3: Chaining of subsystems in a service template	36
Figure 4: Defining subsystem details in a service template.....	39
Figure-5: Typical 3-Tier Network.....	176
Figure-6: Generic Service Template	185
Figure-7: Service template with network template A	185
Figure-8: Service template with network template B.....	186

1 Objective

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

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

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

12 2 Summary of key TOSCA concepts

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

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

29 The TOSCA simple profile assumes a number of base types (node types and relationship types) to be supported
30 by each compliant environment such as a 'Compute' node type, a 'Network' node type or a generic 'Database'
31 node type (see Appendix C). Furthermore, it is envisioned that a large number of additional types for use in
32 service templates will be defined by a community over time. Therefore, template authors in many cases will not
33 have to define types themselves but can simply start writing service templates that use existing types. In addition,
34 the simple profile will provide means for easily customizing existing types, for example by providing a customized
35 'create' script for some software.

36

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

37

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

44 *Example 1 - TOSCA Simple "Hello World"*

```
tosca_definitions_version: tosca_simple_yaml_1_0_0

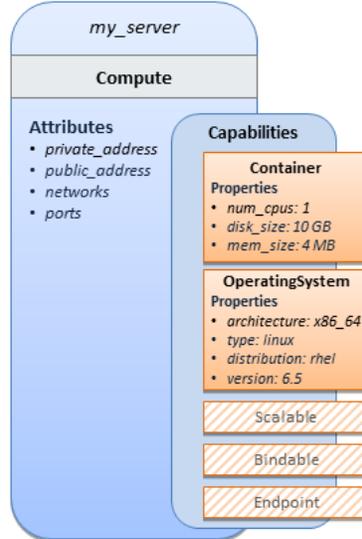
description: Template for deploying a single server with predefined properties.

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: 4 MB
        # Guest Operating System properties
        os:
          properties:
            # host Operating System image properties
            architecture: x86_64
            type: linux
            distribution: rhel
            version: 6.5
```

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

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

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



59
60

61 As you can see, the **Compute** node also has attributes and other built-in capabilities, such as **Bindable** and
 62 **Endpoint**, each with additional properties that will be discussed in other examples later in this document.
 63 Although the Compute node has no direct properties apart from those in its capabilities, other TOSCA node type
 64 definitions may have properties that are part of the node type itself in addition to having Capabilities. TOSCA
 65 orchestration engines are expected to validate all property values provided in a node template against the
 66 property definitions in their respective node type definitions referenced in the service template. The
 67 **tosca_definitions_version** keyname in the TOSCA service template identifies the versioned set of normative
 68 TOSCA type definitions to use for validating those types defined in the TOSCA Simple Profile including the
 69 Compute node type. Specifically, the value **tosca_simple_yaml_1_0_0** indicates Simple Profile v1.0.0 definitions
 70 would be used for validation. Other type definitions may be imported from other service templates using the
 71 **import** keyword discussed later.

72 3.1 Requesting input parameters and providing output

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

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

```
tosca_definitions_version: toscasimpleyaml_1_0_0

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

topology_template:
  inputs:
    cpus:
```

```

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

node_templates:
  my_server:
    type: tosca.nodes.Compute
    capabilities:
      # Host container properties
      host:
        properties:
          # Compute properties
          num_cpus: { get_input: cpus }
          mem_size: 4 MB
          disk_size: 10 GB

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

```

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

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

86

4 TOSCA template for a simple software installation

87

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

89

90

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

```
tosca_definitions_version: tosca_simple_yaml_1_0_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
```

92

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

93

94

95

96

97

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

98

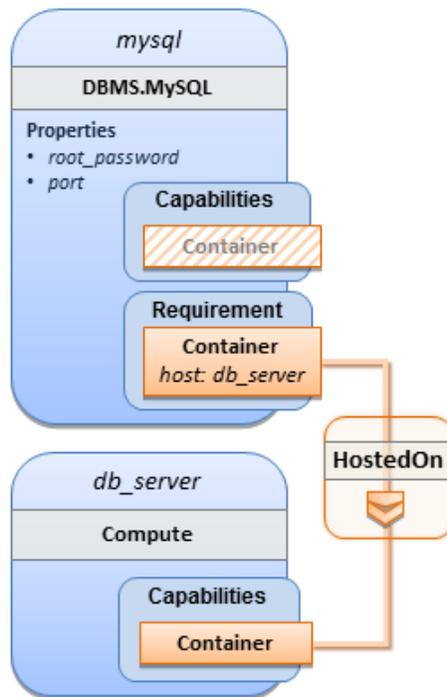
99

100

101

102

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



104

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

114 5 Overriding behavior of predefined node types

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

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

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

```
tosca_definitions_version: tosca_simple_yaml_1_0_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
```

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

128 In other words, operations defined by node types can be thought of as "hooks" into which automation can be
129 injected. Typically, node type implementations provide the automation for those "hooks". However, within a

130 template, custom automation can be injected to run in a hook in the context of the one, specific node template
131 (i.e. without changing the node type).

132

6 TOSCA template for database content deployment

133

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

134

135

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

```
tosca_definitions_version: tosca_simple_yaml_1_0_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:
          implementation: 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
```

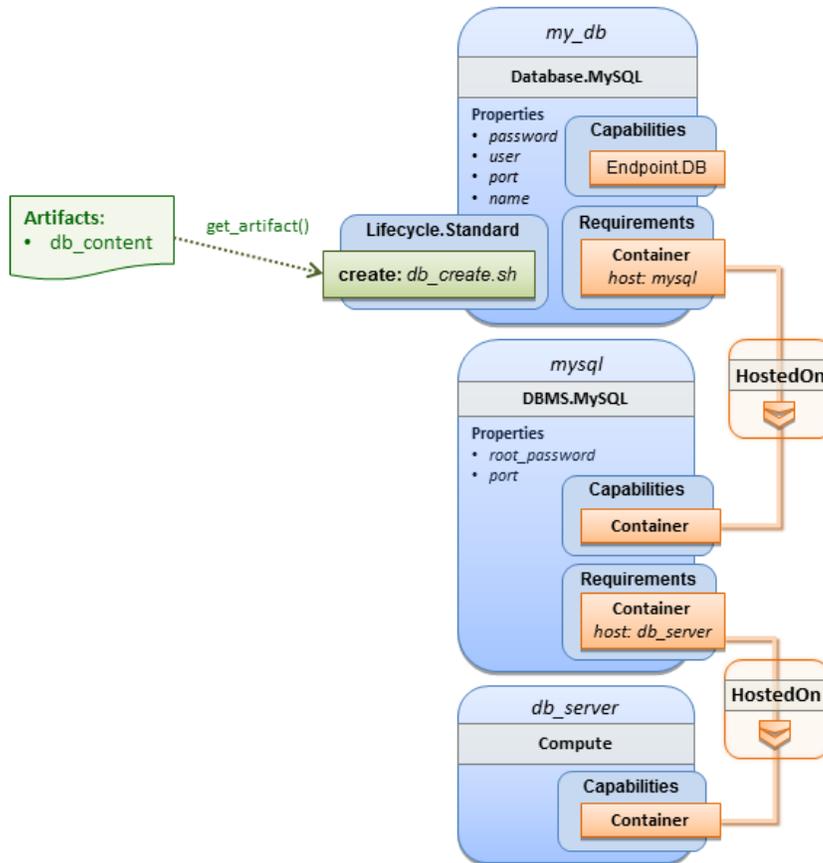
137 In the example above, the **my_db** node template or type **tosca.nodes.Database.MySQL** represents an actual
138 MySQL database instance managed by a MySQL DBMS installation. The **requirements** section of the **my_db**
139 node template expresses that the database it represents is to be hosted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

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

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

153

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



155

156

157

158

159

160

161

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

162

7 TOSCA template for a two-tier application

163

The definition of multi-tier applications in TOSCA is quite similar to the example shown in section 4, with the only difference that multiple software node stacks (i.e., node templates for middleware and application layer components), typically hosted on different servers, are defined and related to each other. The example below defines a web application stack hosted on the **web_server** “compute” resource, and a database software stack similar to the one shown earlier in section 6 hosted on the **db_server** compute resource.

164

165

166

167

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

```
tosca_definitions_version: tosca_simple_yaml_1_0_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: toska.nodes.Compute
  capabilities:
    # omitted here for brevity
```

169 The web application stack consists of the **wordpress**, the **apache** and the **web_server** node templates. The
170 **wordpress** node template represents a custom web application of type
171 **tosca.nodes.WebApplication.WordPress** which is hosted on an Apache web server represented by the **apache**
172 node template. This hosting relationship is expressed via the **host** entry in the **requirements** section of the
173 **wordpress** node template. The **apache** node template, finally, is hosted on the **web_server** compute node.

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

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

187
188

8 Using a custom script to establish a relationship in a template

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

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

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

```
tosca_definitions_version: tosca_simple_yaml_1_0_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 ...
```

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

207

208 The **ConnectsTo** relationship (see C.5.4) also provides a default **Configure** interface with operations that
209 optionally get executed when the orchestrator establishes the relationship. In the above example, the author has
210 provided the custom script **wp_db_configure.sh** to be executed for the operation called
211 **pre_configure_source**. The script file is assumed to be located relative to the referencing service template such
212 as a relative directory within the TOSCA Cloud Service Archive (CSAR) packaging format. This approach allows
213 for conveniently hooking in custom behavior without having to define a completely new derived relationship type.

214

9 Using custom relationship types in a TOSCA template

215

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

217

218

219

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

220

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

```
tosca_definitions_version: tosca_simple_yaml_1_0_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 ...
```

221

In the example above, a special relationship type **my.types.WordpressDbConnection** is specified for

222

establishing the link between the **wordpress** node and the **wordpress_db** node through the use of the

223

relationship (keyword) attribute in the **database** reference. It is assumed, that this special relationship type

224

provides some extra behavior (e.g., an operation with a script) in addition to what a generic “connects to”

225

relationship would provide. The definition of this custom relationship type is shown in the following section.

226 9.1 Definition of a custom relationship type

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

233 *Example 9 - Defining a custom relationship type*

```
tosca_definitions_version: tosca_simple_yaml_1_0_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
```

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

236

10 Defining generic dependencies between nodes in a template

237

238 In some cases it can be necessary to define a generic dependency between two nodes in a template to influence
239 orchestration behavior, i.e. to first have one node processed before another dependent node gets processed. This
240 can be done by using the generic **dependency** requirement which is defined by the [TOSCA Root Node Type](#) and
241 thus gets inherited by all other node types in TOSCA (see section C.7.1).

242

Example 10 - Simple dependency relationship between two nodes

```
tosca_definitions_version: tosca_simple_yaml_1_0_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
```

243 As in previous examples, the relation that one node depends on another node is expressed in the
244 **requirements** section using the built-in requirement named **dependency** that exists for all node types in TOSCA.
245 Even if the creator of the **MyApplication** node type did not define a specific requirement for **SomeService**
246 (similar to the **database** requirement in the example in section 8), the template author who knows that there is
247 a timing dependency and can use the generic **dependency** requirement to express that constraint using the very
248 same syntax as used for all other references.

11 Describing abstract requirements for nodes and capabilities in a TOSCA template

In TOSCA templates, nodes are either:

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

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

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

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

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

11.1 Using a `node_filter` to define hosting infrastructure requirements for a software

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

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

```
tosca_definitions_version: tosca_simple_yaml_1_0_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

```

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

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

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

300 11.2 Using an abstract node template to define infrastructure requirements 301 for software

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

```
tosca_definitions_version: toska_simple_yaml_1_0_0

description: Template with requirements against hosting infrastructure.

topology_template:
  inputs:
    # omitted here for brevity

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

    mysql_compute:
      type: Compute
      capabilities:
        host:
          properties:
            num_cpus: { equal: 2 }
            mem_size: { greater_or_equal: 2 GB }
        os:
          properties:
            architecture: { equal: x86_64 }
            type: linux
            distribution: ubuntu
```

307 As you can see the resulting `mysql_compute` node template looks very much like the “hello world” template as
308 shown in Chapter 3 (where the **Compute** node template was abstract), but this one also allows the TOSCA
309 orchestrator more flexibility when “selecting” a host **Compute** node by providing flexible constraints for
310 properties like `mem_size`.

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

316 11.3 Using a node_filter to define requirements on a database for an 317 application

318 In the same way requirements can be defined on the hosting infrastructure (as shown above) for an application, it
319 is possible to express requirements against application or middleware components such as a database that is not

320 defined in the same template. The provider may then allocate a database by any means, (e.g. using a database-
321 as-a-service solution).

```
tosca_definitions_version: toska_simple_yaml_1_0_0

description: Template with a database requirement.

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 }
```

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

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

331 12 Using node template substitution for model 332 composition

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

338 12.1 Understanding node template instantiation through a TOSCA 339 Orchestrator

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

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

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

352 12.2 Definition of the top-level service template

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

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

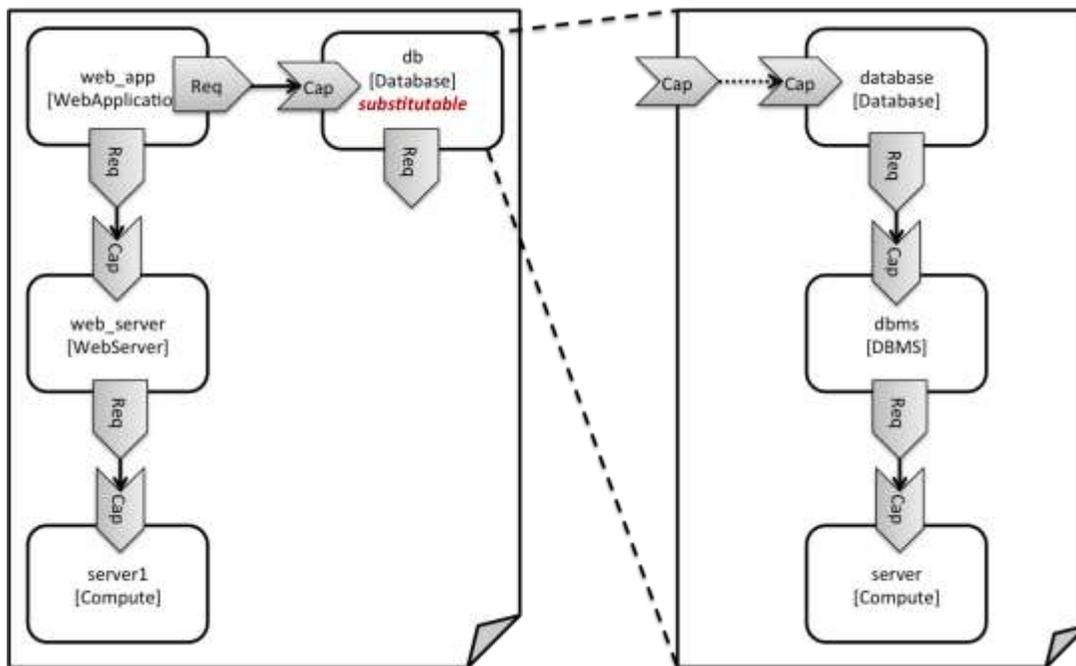


Figure 1: Using template substitution to implement a database tier

361

362

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

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

```

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

```

376 12.3 Definition of the database stack in a service template

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

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

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

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

- 389 1. It defines what node templates, i.e. node templates of which type, can be substituted by the topology
390 template.
- 391 2. It defines how capabilities of the substituted node (or the capabilities defined by the node type of the
392 substituted node template, respectively) are bound to capabilities of node templates defined in the
393 topology template.
- 394 3. It defines how requirements of the substituted node (or the requirements defined by the node type of
395 the substituted node template, respectively) are bound to requirements of node templates defined in
396 the topology template.

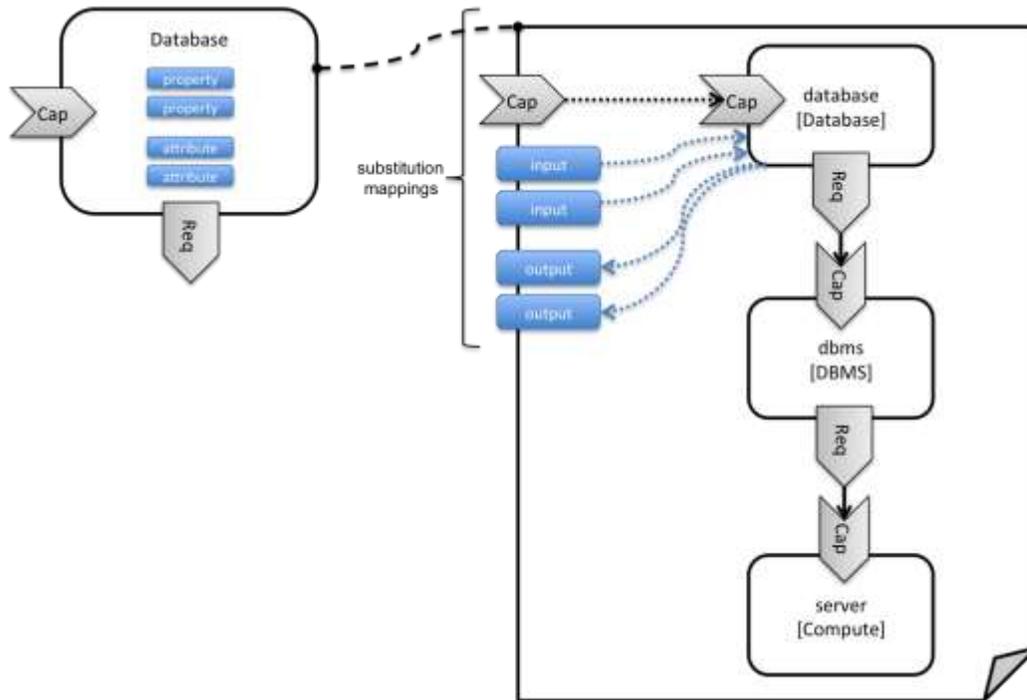


Figure 2: Substitution mappings

397

398

399

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

```

tosca_definitions_version: tosca_simple_yaml_1_0

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

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

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

  node_templates:
    database:
      type: tosca.nodes.Database
      properties:
        user: { get_input: db_user }
        # other properties omitted for brevity
      requirements:
        - host: dbms

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

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

```

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

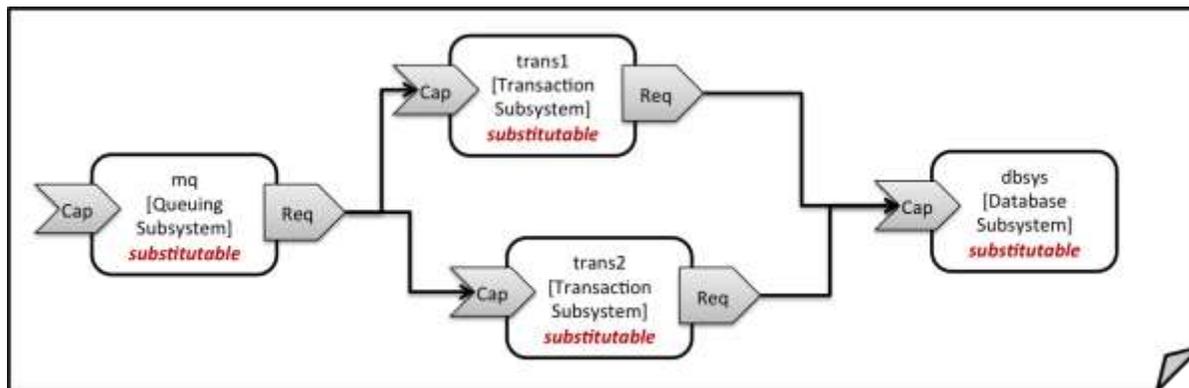
408 Further note that no mappings for properties or attributes of the substituted node are defined. Instead, the inputs
409 and outputs defined by the topology template have to match the properties and attributes of the substituted node.
410 If there are more inputs than the substituted node has properties,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

412 13 Using node template substitution for chaining 413 subsystems

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

422 13.1 Defining the overall subsystem chain

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



432
433 *Figure 3: Chaining of subsystems in a service template*

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

439

```
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
```

440

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

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

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

450 13.2 Defining a subsystem (node) type

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

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

456

```
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
```

457

458 Configuration parameters that shall be allowed for customizing the instantiation of any subsystem are defined as
459 properties of the node type. In the current example, those are the properties **mq_service_ip** and **receiver_port**
460 that had been used in the end-to-end service template in section 13.1.

492 in substitution scenarios. Only the presence of the **substitution_mappings** metadata section in the
493 **topology_template** enables the service template for substitution use cases.
494

```

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 ] }
```

495

14 Grouping node templates

496

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

497

498

499

500

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

501

502

503

Example 11 - Grouping Node Templates with same scaling policy

```
tosca_definitions_version: tosca_simple_yaml_1_0_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:
      members: [ apache, server ]
      policies:
        - my_scaling_policy:
            # Specific policy definitions are considered domain specific and
            # are not included here
```

504

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

505

506

507

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

510 In cases where no explicit binding between software components and their hosting compute resources is defined
511 in a template, but only requirements are defined as has been shown in section 11, a provider could decide to
512 place software components on the same host if their hosting requirements match, or to place them onto different
513 hosts.

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

```
tosca_definitions_version: tosca_simple_yaml_1_0_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: 2 MB }
                      - disk_size: { greater_or_equal: 10 MB }
                - 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_collocation_group:
  members: [ wordpress_server, mysql ]
  policies:
    - my_anti_collocation_policy:
      # Specific policy definitions are considered domain specific and
      # are not included here
```

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

521

15 Using YAML Macros to simplify templates

522

The YAML 1.2 specification allows for defining of [aliases](#) which allow for authoring a block of YAML (or node) once and indicating it is an “anchor” and then referencing it elsewhere in the same document as an “alias”.

523

Effectively, YAML parsers treat this as a “macro” and copy the anchor block’s code to wherever it is referenced.

524

Use of this feature is especially helpful when authoring TOSCA Service Templates where similar definitions and property settings may be repeated multiple times when describing a multi-tier application.

525

526

527

528

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

529

530

531

To accomplish this, the author would describe the reusable properties using a named anchor in the

532

“**dsl_definitions**” section of the TOSCA Service Template and reference the anchor name as an alias in any

533

Compute node templates where these properties may need to be reused. For example:

```
tosca_definitions_version: tosca_simple_yaml_1_0_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: 4096 kB

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
```

534
535

16 Passing information as inputs to Nodes and Relationships

536
537
538
539
540

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

541
542

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

```
node_templates:
  wordpress:
    type: tosca.nodes.WebApplication.WordPress
    requirements:
      ...
      - database_endpoint: mysql_database
    interfaces:
      Standard:
        inputs:
          wp_db_port: { get_property: [ SELF, database_endpoint, port ] }
```

543

16.2 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 ] }
```

544
545
546

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

547 16.3 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
```

548

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

551 16.4 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 ] }
```

552 In this example, the **Standard** lifecycle’s **create** operation exposes / exports an environment variable named
553 “**data_dir**” which will be passed as an input to the **Standard** lifecycle’s **configure** operation.

554

17 Topology Template Model versus Instance Model

555
556
557
558
559
560
561
562
563
564

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

565

18 Using attributes implicitly reflected from properties

566

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

567

568

569

570

571

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

572

573

574

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

575

576

577

578

579

580

```
tosca_definitions_version: toscasimpleyaml_1_0_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

```

581

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

587

588 It should be noted that the actual port value of the **notification_port** attribute may or may not be the value
 589 **8000** as requested on the property; therefore, any node that is dependent on knowing its actual “runtime” value
 590 would use the **get_attribute** function instead of the **get_property** function.

591

Appendix A. TOSCA Simple Profile definitions in YAML

592
593

This section describes all of the YAML 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).

594

A.1 TOSCA namespace and alias

595
596

The following table defines the namespace alias and (target) namespace values that SHALL be used when referencing the TOSCA Simple Profile version 1.0 specification.

Alias	Target Namespace	Specification Description
tosca_simple_yaml_1_0_0	http://docs.oasis-open.org/tosca/ns/simple/yaml/1.0	The TOSCA Simple Profile v1.0 (YAML) target namespace and namespace alias.

597

A.1.1 Rules to avoid namespace collisions

598
599
600
601
602

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.

603
604

The following cases are considered:

605
606
607
608
609
610

- Duplicate property names within same entity (e.g., Node Type, Node Template, Relationship Type, etc.)
- Duplicate requirement names within same entity (e.g., Node Type, Node Template, Relationship Type, etc.)
- Duplicate capability names within same entity (e.g., Node Type, Node Template, Relationship Type, etc.)
- Collisions that occurs from “import” for any Type or Template.
- Collision that occurs from “substitution” of other Node Templates.

611

A.2 Parameter and property types

612
613

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

614

A.2.1 Referenced YAML Types

615
616

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

617
618

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

Valid aliases	Type URI
string	tag:yaml.org,2002:str (default)
integer	tag:yaml.org,2002:int
float	tag:yaml.org,2002:float
boolean	tag:yaml.org,2002:bool (i.e., a value either ‘true’ or ‘false’)

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

619 **A.2.1.1 Notes**

- 620
- The “string” type is the default type when not specified on a parameter or property declaration.
 - While YAML supports further type aliases, such as “str” for “string”, the TOSCA Simple Profile specification promotes the fully expressed alias name for clarity.
- 621
- 622

623 **A.2.2 TOSCA version**

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

Shorthand Name	version
Type Qualified Name	tosca:version

629 **A.2.2.1 Grammar**

630 TOSCA version strings have the following grammar:

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

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

- **major_version**: is a required integer value greater than or equal to 0 (zero)
 - **minor_version**: is a required integer value greater than or equal to 0 (zero).
 - **fix_version**: is a optional integer value greater than or equal to 0 (zero).
 - **qualifier**: is an optional string that indicates a named, pre-release version of the associated code that has been derived from the version of the code identified by the combination **major_version**, **minor_version** and **fix_version** numbers.
 - **build_version**: is an optional integer value greater than or equal to 0 (zero) that can be used to further qualify different build versions of the code that has the same **qualifer_string**.
- 632
- 633
- 634
- 635
- 636
- 637
- 638
- 639

640 **A.2.2.2 Version Comparison**

- When comparing TOSCA versions, all component versions (i.e., major, minor and fix) are compared in sequence from left to right.
 - TOSCA versions that include the optional qualifier are considered older than those without a qualifier.
 - TOSCA versions with the same major, minor, and fix versions and have the same qualifier string, but with different build versions can be compared based upon the build version.
 - Qualifier strings are considered domain-specific. Therefore, this specification makes no recommendation on how to compare TOSCA versions with the same major, minor and fix versions, but with different qualifiers strings and simply considers them different named branches derived from the same code.
- 641
- 642
- 643
- 644
- 645
- 646
- 647
- 648
- 649

650 A.2.2.3 Examples

651 Example of a version with

```
# basic version strings
6.1
2.0.1

# version string with optional qualifier
3.1.0.beta

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

652 A.2.2.4 Notes

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

654 A.2.2.5 Additional Requirements

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

657 A.2.3 TOCSA range type

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

Shorthand Name	range
Type Qualified Name	tosca:range

660 A.2.3.1 Grammar

661 TOSCA range values have the following grammar:

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

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

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

666 A.2.3.2 Keywords:

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

Keyword	Applicable Types	Description
----------------	-------------------------	--------------------

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

668 **A.2.3.3 Examples**

669 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 ]
```

670 **A.2.4 TOSCA list type**

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

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

Shorthand Name	list
Type Qualified Name	tosca:list

676 **A.2.4.1 Grammar**

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

678 **A.2.4.1.1 Square bracket notation**

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

679 **A.2.4.1.2 Bulleted (sequenced) list notation**

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

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

- 681
- **<list_entry_*>**: represents one entry of the list.

682 **A.2.4.2 Declaration Examples**

683 **A.2.4.2.1 List declaration using a simple type**

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

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

686 **A.2.4.2.2 List declaration using a complex type**

687 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
```

688 **A.2.4.3 Definition Examples**

689 These examples show two notation options for defining lists:

- 690
- 691 • A single-line option which is useful for only short lists with simple entries.
 - 692 • A multi-line option where each list entry is on a separate line; this option is typically useful or more readable if there is a large number of entries, or if the entries are complex.

693 **A.2.4.3.1 Square bracket notation**

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

694 **A.2.4.3.2 Bulleted list notation**

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

695 **A.2.5 TOSCA map type**

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

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

Shorthand Name	map
Type Qualified Name	tosca:map

702 **A.2.5.1 Grammar**

703 TOSCA maps are normal YAML dictionaries with following grammar:

704 **A.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>
```

705 **A.2.5.1.2 Multi-line grammar**

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

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

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

709 **A.2.5.2 Declaration Examples**

710 **A.2.5.2.1 Map declaration using a simple type**

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

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

```
- max_length: 128
```

713 **A.2.5.2.2 Map declaration using a complex type**

714 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
```

715 **A.2.5.3 Definition Examples**

716 These examples show two notation options for defining maps:

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

720 **A.2.5.3.1 Single-line notation**

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

721 **A.2.5.3.2 Multi-line notation**

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

722 **A.2.6 TOCSA scalar-unit type**

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

725 **A.2.6.1 Grammar**

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

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

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

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

730 **A.2.6.2 Additional requirements**

- 731
- **Whitespace**: any number of spaces (including zero or none) **SHALL** be allowed between the **scalar** value and the **unit** value.
 - It **SHALL** be considered an error if either the scalar or unit portion is missing on a property or attribute declaration derived from any scalar-unit type.
 - 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]`

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

742 **A.2.6.3 Concrete Types**

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

743

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

- 745
- **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.

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

749 **A.2.6.4 scalar-unit.size**

750 **A.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)

751 **A.2.6.4.2 Examples**

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

752 **A.2.6.4.3 Notes**

- 753
- 754
- 755
- 756
- 757
- 758
- 759
- 760
- The unit values recognized by TOSCA Simple Profile for size-type units are based upon a subset of those defined by GNU at http://www.gnu.org/software/parted/manual/html_node/unit.html, which is a non-normative reference to this specification.
 - TOSCA treats these unit values as case-insensitive (e.g., a value of 'kB', 'KB' or 'kb' would be equivalent), but it is considered best practice to use the case of these units as prescribed by GNU.
 - Some Cloud providers may not support byte-level granularity for storage size allocations. In those cases, these values could be treated as desired sizes and actual allocations would be based upon individual provider capabilities.

761 **A.2.6.5 scalar-unit.time**

762 **A.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

763 **A.2.6.5.2 Examples**

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

764 **A.2.6.5.3 Notes**

- 765
- 766
- 767
- 768
- 769
- The unit values recognized by TOSCA Simple Profile for time-type units are based upon a subset of those defined by International System of Units whose recognized abbreviations are defined within the following reference:
 - <http://www.ewh.ieee.org/soc/ias/pub-dept/abbreviation.pdf>
 - This document is a non-normative reference to this specification and intended for publications or grammars enabled for Latin characters which are not accessible in typical programming languages

770 A.2.6.6 scalar-unit.frequency

771 A.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.

772 A.2.6.6.2 Examples

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

773 A.2.6.6.3 Notes

- 774 • The value for Hertz (Hz) is the International Standard Unit (ISU) as described by the Bureau International
775 des Poids et Mesures (BIPM) in the “*SI Brochure: The International System of Units (SI) [8th edition,*
776 *2006; updated in 2014]*”, <http://www.bipm.org/en/publications/si-brochure/>

777 A.3 Normative values

778 A.3.1 Node States

779 As components (i.e., nodes) of TOSCA applications are deployed, instantiated and orchestrated over their
780 lifecycle using normative lifecycle operations (see section C.6 for normative lifecycle definitions) it is important
781 define normative values for communicating the states of these components normatively between orchestration
782 and workflow engines and any managers of these applications.

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

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

Node State		
Value	Transitional	Description
starting	yes	Node is transitioning from configured state to started state.
started	no	Node is started.
stopping	yes	Node is transitioning from its current state to a configured state.
deleting	yes	Node is transitioning from its current state to one where it is deleted and its state is no longer tracked by the instance model.
error	no	Node is in an error state.

785 A.3.2 Directives

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

787 A.3.3 Network Name aliases

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

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

791 A.3.3.1 Usage

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

795 A.4 TOSCA Metamodel

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

798 A.4.1 Required Keynames

799 The TOSCA metamodel includes complex types (e.g., Node Types, Relationship Types, Capability Types, Data
 800 Types, etc.) each of which include their own list of reserved keynames that are sometimes marked as **required**.

801 These types may be used to derive other types. These derived types (e.g., child types) do not have to provide
802 required keynames as long as they have been specified in the type they have been derived from (i.e., their parent
803 type).

804 **A.5 Reusable modeling definitions**

805 **A.5.1 Description definition**

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

808 **A.5.1.1 Keyname**

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

```
description
```

810 **A.5.1.2 Grammar**

811 Description definitions have the following grammar:

```
description: <string>
```

812 **A.5.1.3 Examples**

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

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

815 The YAML “folded” style may also be used for multi-line descriptions which “folds” line breaks as space
816 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.
```

817 **A.5.1.4 Notes**

- 818
- 819 • Use of “folded” style is discouraged for the YAML string type apart from when used with the **description** keyname.

820 **A.5.2 Constraint clause**

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

824 **A.5.2.1 Operator keynames**

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

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

826 **A.5.2.1.1 Comparable value types**

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

829 **A.5.2.2 Additional Requirements**

- 830
- 831 If no operator is present for a simple scalar-value on a constraint clause, it **SHALL** be interpreted as being equivalent to having the “**equal**” operator provided; however, the “**equal**” operator may be used for clarity when expressing a constraint clause.
 - 832
 - 833 The “**length**” operator **SHALL** be interpreted mean “size” for set types (i.e., list, map, etc.).

834 **A.5.2.3 Grammar**

835 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>
```

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

- 837 • **operator**: represents a required operator from the specified list shown above (see section A.5.2.1
- 838 “Operator keynames”).
- 839 • **scalar_value**, **scalar_value_***: represents a required scalar (or atomic quantity) that can hold only
- 840 one value at a time. This will be a value of a primitive type, such as an integer or string that is allowed
- 841 by this specification.
- 842 • **value_***: represents a required value of the operator that is not limited to scalars.
- 843 • **regular_expression_value**: represents a regular expression (string) value.

844 A.5.2.4 Examples

845 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
```

846 A.5.2.5 Additional Requirements

- 847
- Values provided by the operands (i.e., values and scalar values) **SHALL** be type-compatible with their associated operations.
 - Future drafts of this specification will detail the use of regular expressions and reference an appropriate standardized grammar.
- 848
- 849
- 850

851 A.5.3 Property Filter definition

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

854 A.5.3.1 Grammar

855 Property filter definitions have one of the following grammars:

856 A.5.3.1.1 Short notation:

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

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

858 A.5.3.1.2 Extended notation:

859 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>
```

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

- 861
- **property_name**: represents the name of property that would be used to select a property definition with the same name (**property_name**) on a TOSCA entity (e.g., a Node Type, Node Template, Capability Type, etc.).
 - **property_constraint_clause_***: represents constraint clause(s) that would be used to filter entities based upon the named property's value(s).
- 862
- 863
- 864
- 865

866 A.5.3.2 Additional Requirements

- 867
- Property constraint clauses must be type compatible with the property definitions (of the same name) as defined on the target TOSCA entity that the clause would be applied against.
- 868

869 A.5.4 Node Filter definition

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

872 A.5.4.1 Keynames

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

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

874 A.5.4.2 Additional filtering on named Capability properties

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

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

876 A.5.4.3 Grammar

877 Node filter definitions have one of the following grammars:

```
<filter_name>:  
  properties:  
    - <property_filter_def_1>  
    - ...  
    - <property_filter_def_n>  
  capabilities:  
    - <capability_name_or_type_1>:  
      properties:  
        - <cap_1_property_filter_def_1>  
        - ...  
        - <cap_m_property_filter_def_n>  
    - ...  
    - <capability_name_or_type_n>:  
      properties:  
        - <cap_1_property_filter_def_1>  
        - ...  
        - <cap_m_property_filter_def_n>
```

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

- 879
- **property_filter_def_***: represents a property filter definition that would be used to select (filter) matching TOSCA entities (e.g., Node Template, Node Type, Capability Types, etc.) based upon their property definitions' values.
 - **property_constraint_clause_***: represents constraint clause(s) that would be used to filter entities based upon property values.
 - **capability_name_or_type_***: represents the type or name of a capability that would be used to select (filter) matching TOSCA entities based upon their existence.
 - **cap_*_property_def_***: represents a property filter definition that would be used to select (filter) matching TOSCA entities (e.g., Node Template, Node Type, Capability Types, etc.) based upon their capabilities' property definitions' values.
- 880
881
882
883
884
885
886
887
888

889 A.5.4.4 Additional requirements

- TOSCA orchestrators **SHALL** search for matching capabilities listed on a target filter by assuming the capability name is first a symbolic name and secondly it is a type name (in order to avoid namespace collisions).
- 890
891
892

893 A.5.4.5 Example

894 The following example is a filter that would be used to select a TOSCA **Compute** node based upon the values of
895 its defined capabilities. Specifically, this filter would select Compute nodes that supported a specific range of
896 CPUs (i.e., **num_cpus** value between 1 and 4) and memory size (i.e., **mem_size** of 2 or greater) from its declared
897 "host" capability. In this example, the author also wants the Compute node to support an encryption capability of
898 type **mytypes.capabilities.compute.encryption** which has properties that support a specific (AES)
899 encryption **algorithm** and **keylength** (128).

900

```
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: 2 MB }
          # and should also support this type of encryption and properties:
          - mytypes.capabilities.compute.encryption:
              properties:
                - algorithm: { equal: aes }
                - keylength: { valid_values: [ 128, 256 ] }
```

901 A.5.5 Artifact definition

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

904 A.5.5.1 Keynames

905 The following is the list of recognized keynames recognized for a TOSCA artifact definition:

Keyname	Required	Type	Description
type	yes	string	The required artifact type for the artifact definition.
implementation	no	string	The optional 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 implementation 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.

906 A.5.5.2 Grammar

907 Artifact definitions have one of the following grammars:

908 A.5.5.2.1 Short notation

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

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

911 A.5.5.2.2 Extended notation:

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

```
<artifact_name>:  
  description: <artifact_description>  
  type: <artifact_type_name>  
  implementation: <artifact_file_URI>  
  deploy_path: <file_deployment_path>
```

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

- 915 • **artifact_name**: represents the required symbolic name of the artifact as a **string**.
- 916 • **artifact_description**: represents the optional **description** for the artifact.
- 917 • **artifact_type_name**: represents the required **artifact type** the artifact definition is based upon.
- 918 • **artifact_file_URI**: represents the required **URI string (relative or absolute) which can be used to**
919 **locate the artifact's file.**
- 920 • **file_deployment_path**: represents the optional path the **artifact_file_URI** would be copied into
921 within the target node's container.

922 A.5.5.3 Example

923 The following represents an artifact definition:

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

924 A.5.6 Repository definition

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

927 A.5.6.1 Keynames

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

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

929 A.5.6.2 Grammar

930 Repository definitions have one the following grammars:

931 A.5.6.2.1 Single-line grammar (no credential):

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

932

933 A.5.6.2.2 Multi-line grammar

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

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

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

940 A.5.6.3 Additional Requirements

- 941 • None

942 A.5.6.4 Example

943 The following represents a repository definition:

```
repositories:  
  my_code_repo:  
    description: My project's code repository in GitHub
```

url: <https://github.com/my-project/>

944 A.5.7 Property definition

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

950 A.5.7.1.1 Attribute and Property reflection

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

954 A.5.7.2 Keynames

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

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

956 A.5.7.3 Status values

957 The following property status values are supported:

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

958 A.5.7.4 Grammar

959 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>

```

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

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

977 A.5.7.5 Additional Requirements

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

985 A.5.7.6 Notes

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

990 **A.5.7.7 Example**

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

```
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 ]
```

992 **A.5.8 Property assignment**

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

995 **A.5.8.1 Keynames**

996 The TOSCA property assignment has no keynames.

997 **A.5.8.2 Grammar**

998 Property assignments have the following grammar:

999 **A.5.8.2.1 Short notation:**

1000 The following single-line grammar may be used when a simple value assignment is needed:

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

1001 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1002 • **property_name:** represents the name of a property that would be used to select a property definition
1003 with the same name within on a TOSCA entity (e.g., Node Template, Relationship Template, etc.) which
1004 is declared in its declared type (e.g., a Node Type, Node Template, Capability Type, etc.).
- 1005 • **property_value, property_value_expression:** represent the type-compatible value to assign to
1006 the named property. Property values may be provided as the result from the evaluation of an
1007 expression or a function.

1008 **A.5.9 Attribute definition**

1009 An attribute definition defines a named, typed value that can be associated with an entity defined in this
1010 specification (e.g., a Node, Relationship or Capability Type). Specifically, it is used to expose the “actual state” of
1011 some property of a TOSCA entity after it has been deployed and instantiated (as set by the TOSCA orchestrator).
1012 Attribute values can be retrieved via the **get_attribute** function from the instance model and used as values to
1013 other entities within TOSCA Service Templates.

1014 **A.5.9.1.1 Attribute and Property reflection**

1015 TOSCA orchestrators automatically create [Attribute definitions](#) for any [Property definitions](#) declared on the same
1016 TOSCA entity (e.g., nodes, node capabilities and relationships) in order to make accessible the actual (i.e., the
1017 current state) value from the running instance of the entity.

1018 A.5.9.2 Keynames

1019 The following is the list of recognized keynames recognized for a TOSCA attribute definition:

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

1020 A.5.9.3 Grammar

1021 Attribute definitions have the following grammar:

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

1022 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1023 • **attribute_name**: represents the required symbolic name of the attribute as a [string](#).
- 1024 • **attribute_type**: represents the required data type of the attribute.
- 1025 • **attribute_description**: represents the optional [description](#) of the attribute.
- 1026 • **default_value**: contains a type-compatible value that may be used as a default if not provided by
1027 another means.
- 1028 • **status_value**: contains a value indicating the attribute's status relative to the specification version (e.g.,
1029 supported, deprecated, etc.). Supported [status values](#) for this keyname are defined under [Property](#)
1030 [definition](#).

1031 A.5.9.4 Additional Requirements

- 1032 • In addition to any explicitly defined attributes on a TOSCA entity (e.g., Node Type, RelationshipType,
1033 etc.), implementations of the TOSCA Simple Profile **MUST** automatically reflect (i.e., make available) any
1034 property defined on an entity as an attribute of the entity with the same name as the property.
- 1035 • Values for the default keyname **MUST** be derived or calculated from other attribute or operation output
1036 values (that reflect the actual state of the instance of the corresponding resource) and not hard-coded
1037 or derived from a property settings or inputs (i.e., desired state).

1038 A.5.9.5 Notes

- 1039 • Attribute definitions are very similar to [Property definitions](#); however, properties of entities reflect an
1040 input that carries the template author's requested or desired value (i.e., desired state) which the

1041 orchestrator (attempts to) use when instantiating the entity whereas attributes reflect the actual value
1042 (i.e., actual state) that provides the actual instantiated value.
1043 ○ For example, a property can be used to request the IP address of a node using a property
1044 (setting); however, the actual IP address after the node is instantiated may be different and
1045 made available by an attribute.

1046 **A.5.9.6 Example**

1047 The following represents a required attribute definition:

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

1048 **A.5.10 Attribute assignment**

1049 This section defines the grammar for assigning values to named attributes within TOSCA Node and Relationship
1050 templates which are defined in their corresponding named types.

1051 **A.5.10.1 Keynames**

1052 The TOSCA attribute assignment has no keynames.

1053 **A.5.10.2 Grammar**

1054 Attribute assignments have the following grammar:

1055 **A.5.10.2.1 Short notation:**

1056 The following single-line grammar may be used when a simple value assignment is needed:

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

1057 **A.5.10.2.2 Extended notation:**

1058 The following multi-line grammar may be used when a value assignment requires keys in addition to a simple
1059 value assignment:

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

1060 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1061 • **attribute_name**: represents the name of an attribute that would be used to select an attribute
1062 definition with the same name within on a TOSCA entity (e.g., Node Template, Relationship Template,
1063 etc.) which is declared (or reflected from a Property definition) in its declared type (e.g., a Node Type,
1064 Node Template, Capability Type, etc.).
- 1065 • **attribute_value**, **attribute_value_expression**: represent the type-compatible value to assign to
1066 the named attribute. Attribute values may be provided as the result from the evaluation of an
1067 expression or a function.
- 1068 • **attribute_description**: represents the optional [description](#) of the attribute.

1069 **A.5.10.3 Additional requirements**

- 1070 • Attribute values **MAY** be provided by the underlying implementation at runtime when requested by the
1071 `get_attribute` function or it **MAY** be provided through the evaluation of expressions and/or functions
1072 that derive the values from other TOSCA attributes (also at runtime).

1073 **A.5.11 Operation definition**

1074 An operation definition defines a named function or procedure that can be bound to an implementation artifact
1075 (e.g., a script).

1076 **A.5.11.1 Keynames**

1077 The following is the list of recognized keynames recognized for a TOSCA operation definition:

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

1078 The following is the list of recognized keynames to be used with the **implementation** keyname within a TOSCA
1079 operation definition:

Keyname	Required	Type	Description
primary	no	string	The optional implementation artifact name (e.g., the primary script file name within a TOSCA CSAR file).
dependencies	no	list of string	The optional 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).

1080 **A.5.11.2 Grammar**

1081 Operation definitions have the following grammars:

1082 **A.5.11.2.1 Short notation**

1083 The following single-line grammar may be used when only an operation's implementation artifact is needed:

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

1084 **A.5.11.2.2 Extended notation for use in Type definitions**

1085 The following multi-line grammar may be used in Node or Relationship Type definitions when additional
1086 information about the operation is needed:

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

1087 **A.5.11.2.3 Extended notation for use in Template definitions**

1088 The following multi-line grammar may be used in Node or Relationship Template definitions when additional
1089 information about the operation is needed:

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

1090 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1091 • **operation_name**: represents the required symbolic name of the operation as a [string](#).
- 1092 • **operation_description**: represents the optional [description](#) string for the corresponding
1093 **operation_name**.
- 1094 • **implementation_artifact_name**: represents the optional name ([string](#)) of an implementation artifact
1095 definition (defined elsewhere), or the direct name of an implementation artifact's relative filename (e.g.,
1096 a service template-relative, path-inclusive filename or absolute file location using a URL).
- 1097 • **property_definitions**: represents the optional list of [property definitions](#) which the TOSCA
1098 orchestrator would make available (i.e., or pass) to the corresponding implementation artifact during its
1099 execution.
- 1100 • **property_assignments**: represents the optional list of [property assignments](#) for passing parameters to
1101 Node or Relationship Template operations providing values for properties defined in their respective
1102 type definitions.
- 1103 • **list_of_dependent_artifact_names**: represents the optional list of one or more dependent or
1104 secondary implementation artifact name which are referenced by the primary implementation artifact.

1105 **A.5.11.3 Additional requirements**

- 1106 • The default sub-classing behavior for implementations of operations SHALL be override. That is,
1107 implementation artifacts assigned in subclasses override any defined in its parent class.
- 1108 • Template authors may provide property assignments on operation inputs on templates that do not
1109 necessarily have a property definition defined in its corresponding type.
- 1110 • Implementation artifact file names (e.g., script filenames) may include file directory path names that are
1111 relative to the TOSCA service template file itself when packaged within a TOSCA Cloud Service ARchive
1112 (CSAR) file.

1113 A.5.11.4 Examples

1114 A.5.11.4.1 Single-line implementation example

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

1115 A.5.11.4.2 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
```

1116 A.5.12 Interface definition

1117 An interface definition defines a named interface that can be associated with a Node or Relationship Type

1118 A.5.12.1 Keynames

1119 The following is the list of recognized keynames recognized for a TOSCA interface definition:

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

1120 A.5.12.2 Grammar

1121 Interface definitions have the following grammar:

1122 A.5.12.2.1 Extended notation for use in Type definitions

1123 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>

```

1124 **A.5.12.2.2 Extended notation for use in Template definitions**

1125 The following multi-line grammar may be used in Node or Relationship Type definitions:

```

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

```

1126 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1127 • **interface_definition_name**: represents the required symbolic name of the interface as a [string](#).
- 1128 • **interface_type_name**: represents the required name of the Interface Type for the interface
- 1129 **definition**.
- 1130 • **property_definitions**: represents the optional list of [property definitions](#) (i.e., parameters) which the
- 1131 TOSCA orchestrator would make available (i.e., or pass) to all defined operations.
- 1132 - *This means these properties and their values would be accessible to the implementation artifacts*
- 1133 *(e.g., scripts) associated to each operation during their execution.*
- 1134 • **property_assignments**: represents the optional list of [property assignments](#) for passing parameters to
- 1135 Node or Relationship Template operations providing values for properties defined in their respective
- 1136 type definitions.
- 1137 • **operation_definitions**: represents the required name of one or more [operation definitions](#).

1138 **A.6 Type-specific definitions**

1139 **A.6.1 Capability definition**

1140 A capability definition defines a named, typed set of data that can be associated with Node Type or Node

1141 Template to describe a transparent capability or feature of the software component the node describes.

1142 **A.6.1.1 Keynames**

1143 The following is the list of recognized keynames for a TOSCA capability definition:

Keyname	Required	Type	Constraints	Description
type	yes	string	N/A	The required name of the Capability Type the capability definition is based upon.
description	no	description	N/A	The optional description of the Capability definition.
properties	no	list of property definitions	N/A	An optional list of property definitions for the Capability definition.
attributes	no	list of attribute definitions	N/A	An optional list of attribute definitions for the Capability definition.

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

1144 A.6.1.2 Grammar

1145 Capability definitions have one of the following grammars:

1146 A.6.1.2.1 Short notation

1147 The following grammar may be used when only a list of capability definition names needs to be declared:

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

1148 A.6.1.2.2 Extended notation

1149 The following multi-line grammar may be used when additional information on the capability definition is
1150 needed:

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

1151 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1152 • **capability_definition_name**: represents the symbolic name of the capability as a [string](#).
- 1153 • **capability_type**: represents the required name of a [capability type](#) the capability definition is based
1154 upon.
- 1155 • **capability_description**: represents the optional [description](#) of the capability definition.
- 1156 • **property_definitions**: represents the optional list of [property definitions](#) for the capability definition.
- 1157 • **attribute_definitions**: represents the optional list of [attribute definitions](#) for the capability
1158 definition.
- 1159 • **node_type_names**: represents the optional list of one or more names of [Node Types](#) that the Capability
1160 definition supports as valid sources for a successful relationship to be established to itself.

1161 A.6.1.3 Examples

1162 The following examples show capability definitions in both simple and full forms:

1163 **A.6.1.3.1 Simple notation example**

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

1164 **A.6.1.3.2 Full notation example**

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

1165 **A.6.1.4 Additional requirements**

- 1166 • Any Node Type (names) provides as values for the **valid_source_types** keyname SHALL be type-
1167 compatible (i.e., derived from the same parent Node Type) with any Node Types defined using the same
1168 keyname in the parent Capability Type.
- 1169 • Capability symbolic names SHALL be unique; it is an error if a capability name is found to occur more
1170 than once.

1171 **A.6.1.5 Notes**

- 1172 • The Capability Type, in this example **MyCapabilityTypeName**, would be defined elsewhere and
1173 have an integer property named **limit**.
- 1174 • This definition directly maps to the **CapabilitiesDefinition** of the Node Type entity as defined in the
1175 [TOSCA v1.0 specification](#).

1176 **A.6.2 Requirement definition**

1177 The Requirement definition describes a named requirement (dependencies) of a TOSCA Node Type or Node
1178 template which needs to be fulfilled by a matching Capability definition declared by another TOSCA modelable
1179 entity. The requirement definition may itself include the specific name of the fulfilling entity (explicitly) or provide
1180 an abstract type, along with additional filtering characteristics, that a TOSCA orchestrator can use to fulfil the
1181 capability at runtime (implicitly).

1182 **A.6.2.1 Keynames**

1183 The following is the list of recognized keynames for a TOSCA requirement definition:

Keyname	Required	Type	Constraints	Description
capability	yes	string	N/A	The required reserved keyname used that can be used to provide the name of a valid Capability Type that can fulfil the requirement.
node	no	string	N/A	The optional reserved keyname used to provide the name of a valid Node Type that contains the capability definition that can be used to fulfil the requirement.
relationship	no	string	N/A	The optional reserved keyname used to provide the name of a valid Relationship Type to construct when fulfilling the requirement.

Keyname	Required	Type	Constraints	Description
occurrences	no	range of integer	implied default of [1,1]	The optional minimum and maximum occurrences for the requirement. Note: the keyword UNBOUNDED is also supported to represent any positive integer.

1184 **A.6.2.1.1 Additional Keynames for multi-line relationship grammar**

1185 The Requirement definition contains the Relationship Type information needed by TOSCA Orchestrators to
 1186 construct relationships to other TOSCA nodes with matching capabilities; however, it is sometimes recognized
 1187 that additional properties may need to be passed to the relationship (perhaps for configuration). In these cases,
 1188 additional grammar is provided so that the Node Type may declare additional Property definitions to be used as
 1189 inputs to the Relationship Type's declared interfaces (or specific operations of those interfaces).

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

1190 **A.6.2.2 Grammar**

1191 Requirement definitions have one of the following grammars:

1192 **A.6.2.2.1 Simple grammar (Capability Type only)**

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

1193 **A.6.2.2.2 Extended grammar (with Node and Relationship Types)**

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

1194 **A.6.2.2.3 Extended grammar for declaring Property Definitions on the relationship's Interfaces**

1195 The following additional multi-line grammar is provided for the relationship keyname in order to declare new
 1196 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>
```

1197 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1198 • **requirement_name**: represents the required symbolic name of the requirement definition as a [string](#).
- 1199 • **capability_type_name**: represents the required name of a Capability type that can be used to fulfil the
1200 requirement.
- 1201 • **node_type_name**: represents the optional name of a TOSCA Node Type that contains the Capability
1202 Type definition the requirement can be fulfilled by.
- 1203 • **relationship_type_name**: represents the optional name of a [Relationship Type](#) to be used to construct
1204 a relationship between this requirement definition (i.e., in the source node) to a matching capability
1205 definition (in a target node).
- 1206 • **min_occurrences, max_occurrences**: represents the optional minimum and maximum occurrences of
1207 the requirement (i.e., its cardinality).
- 1208 • **interface_definitions**: represents one or more already declared interfaces in the Relationship Type
1209 (as declared on the **type** keyname) allowing for the declaration of new Property definition for these
1210 interfaces or for specific Operation definitions of these interfaces.

1211 **A.6.2.3 Additional Requirements**

- 1212 • Requirement symbolic names SHALL be unique; it is an error if a requirement name is found to occur
1213 more than once.
- 1214 • If the **occurrences** keyname is not present, then the occurrence of the requirement **SHALL** be one and
1215 only one; that is a default declaration as follows would be assumed:
1216 ◦ **occurrences**: [1,1]

1217 **A.6.2.4 Notes**

- 1218 • This element directly maps to the **RequirementsDefinition** of the Node Type entity as defined in the
1219 [TOSCA v1.0 specification](#).
- 1220 • The requirement symbolic name is used for identification of the requirement definition only and not
1221 relied upon for establishing any relationships in the topology.

1222 **A.6.2.5 Requirement Type definition is a tuple**

1223 A requirement definition allows type designers to govern which types are allowed (valid) for fulfillment using three
1224 levels of specificity with only the Capability Type being required.

- 1225 1. Node Type (optional)
- 1226 2. Relationship Type (optional)
- 1227 3. Capability Type (required)

1228 The first level allows selection, as shown in both the simple or complex grammar, simply providing the node's
1229 type using the **node** keyname. The second level allows specification of the relationship type to use when
1230 connecting the requirement to the capability using the **relationship** keyname. Finally, the specific named
1231 capability type on the target node is provided using the **capability** keyname.

1232 **A.6.2.5.1 Property filter**

1233 In addition to the node, relationship and capability types, a filter, with the keyname **node_filter**, may be
1234 provided to constrain the allowed set of potential target nodes based upon their properties and their capabilities'
1235 properties. This allows TOSCA orchestrators to help find the "best fit" when selecting among multiple potential
1236 target nodes for the expressed requirements.

1237 A.6.3 Artifact Type

1238 An Artifact Type is a reusable entity that defines the type of one or more files which Node Types or Node
1239 Templates can have dependent relationships and used during operations such as during installation or
1240 deployment.

1241 A.6.3.1 Keynames

1242 The following is the list of recognized keynames for a TOSCA Artifact Type definition:

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

1243 A.6.3.2 Grammar

1244 Artifact Types have following grammar:

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

1245 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1246 • **artifact_type_name**: represents the name of the Artifact Type being declared as a **string**.
- 1247 • **parent_artifact_type_name**: represents the **name** of the **Artifact Type** this Artifact Type definition
1248 derives from (i.e., its “parent” type).
- 1249 • **artifact_description**: represents the optional **description** string for the Artifact Type.
- 1250 • **mime_type_string**: represents the optional Multipurpose Internet Mail Extensions (MIME) standard
1251 string value that describes the file contents for this type of Artifact Type as a **string**.
- 1252 • **file_extensions**: represents the optional list of one or more recognized file extensions for this type of
1253 artifact type as **strings**.
- 1254 • **property_definitions**: represents the optional list of **property definitions** for the artifact type.

1255 A.6.3.3 Examples

```
my_artifact_type:  
  description: Java Archive artifact type  
  derived_from: toska.artifact.Root  
  mime_type: application/java-archive
```

```
file_ext: [ jar ]
```

1256 A.6.4 Interface Type

1257 An Interface Type is a reusable entity that describes a set of operations that can be used to interact with or
1258 manage a node or relationship in a TOSCA topology.

1259 A.6.4.1 Keynames

1260 The following is the list of recognized keynames for a TOSCA Interface Type definition:

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

1261 A.6.4.2 Grammar

1262 Interface Types have following grammar:

```
<interface_type_name>:  
  inputs:  
    <property_definitions>  
  <operation_definitions>
```

1263 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1264 • **interface_type_name**: represents the required name of the interface as a [string](#).
- 1265 • **property_definitions**: represents the optional list of [property definitions](#) (i.e., parameters) which the
1266 TOSCA orchestrator would make available (i.e., or pass) to all implementation artifacts for operations
1267 declared on the interface during their execution.
- 1268 • **operation_definitions**: represents the required list of one or more [operation definitions](#).

1269 A.6.4.3 Example

1270 The following example shows a custom interface used to define multiple configure operations.

```
mycompany.mytypes.myinterfaces.MyConfigure:  
  inputs:  
    mode:  
      type: string  
  pre_configure_service:  
    description: pre-configure operation for my service  
  post_configure_service:  
    description: post-configure operation for my service
```

1271 A.6.4.4 Additional Requirements

- 1272 • Interface Types **MUST NOT** include any implementations for defined operations; that is, the
1273 implementation keyname is invalid.

1274 A.6.4.5 Notes

- 1275 • The TOSCA Simple Profile specification does not yet provide a means to derive or extend an Interface
1276 Type from another Interface Type.

1277 A.6.5 Data Type

1278 A Data Type definition defines the schema for new named datatypes in TOSCA.

1279 A.6.5.1 Keynames

1280 The following is the list of recognized keynames for a TOSCA Data Type definition:

Keyname	Required	Type	Description
derived_from	no	string	The optional key used when a datatype is derived from an existing TOSCA Data Type.
description	no	description	The optional description for the Data Type.
constraints	no	list of constraint clauses	The optional list of <i>sequenced</i> constraint clauses for the Data Type.
properties	no	list of property definitions	The optional list property definitions that comprise the schema for a complex Data Type in TOSCA.

1281 A.6.5.2 Grammar

1282 Data Types have the following grammar:

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

1283 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1284 • **data_type_name**: represents the required symbolic name of the datatype as a [string](#).
- 1285 • **datatype_description**: represents the optional [description](#) for the datatype.
- 1286 • **existing_type_name**: represents the optional name of a valid TOSCA type this new datatype would
1287 derive from.
- 1288 • **type_constraints**: represents the optional *sequenced* list of one or more type-compatible [constraint](#)
1289 [clauses](#) that restrict the datatype.
- 1290 • **property_definitions**: represents the optional list of one or more [property definitions](#) that provide
1291 the schema for the datatype.

1292 A.6.5.3 Additional Requirements

- 1293 • A valid datatype definition **MUST** have either a valid **derived_from** declaration or at least one valid
1294 property definition.
- 1295 • Any **constraint** clauses **SHALL** be type-compatible with the type declared by the **derived_from**
1296 keyname.

- If a **properties** keyname is provided, it **SHALL** contain one or more valid property definitions.

1298 **A.6.5.4 Examples**

1299 The following example represents a datatype definition based upon an existing string type:

1300 **A.6.5.4.1 Defining a complex datatype**

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

1301 **A.6.5.4.2 Defining a datatype derived from an existing datatype**

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

1302 **A.6.6 Capability Type**

1303 A Capability Type is a reusable entity that describes a kind of capability that a Node Type can declare to expose.
 1304 Requirements (implicit or explicit) that are declared as part of one node can be matched to (i.e., fulfilled by) the
 1305 Capabilities declared by another node.

1306 **A.6.6.1 Keynames**

1307 The following is the list of recognized keynames for a TOSCA Capability Type definition:

Keyname	Required	Type	Description
derived_from	no	string	An optional parent capability type name this new Capability Type derives from.
description	no	description	An optional description for the Capability Type.
properties	no	list of property definitions	An optional list of property definitions for the Capability Type.

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

1308 A.6.6.2 Grammar

1309 Capability Types have following grammar:

```

<capability_type_name>:
  derived_from: <parent_capability_type_name>
  description: <capability_description>
  properties:
    <property_definitions>
  attributes:
    <attribute_definitions>
  valid_source_types: [ <node_type_names> ]

```

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

- 1311 • **capability_type_name**: represents the required name of the Capability Type being declared as a
- 1312 [string](#).
- 1313 • **parent_capability_type_name**: represents the name of the [Capability Type](#) this Capability Type
- 1314 definition derives from (i.e., its “parent” type).
- 1315 • **capability_description**: represents the optional [description](#) string for the corresponding
- 1316 **capability_type_name**.
- 1317 • **property_definitions**: represents an optional list of [property definitions](#) that the Capability type
- 1318 exports.
- 1319 • **attribute_definitions**: represents the optional list of [attribute definitions](#) for the Capability Type.
- 1320 • **node_type_names**: represents the optional list of one or more names of [Node Types](#) that the Capability
- 1321 Type supports as valid sources for a successful relationship to be established to itself.

1322 A.6.6.3 Example

```

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

```

1323 **A.6.7 Requirement Type**

1324 A Requirement Type is a reusable entity that describes a kind of requirement that a Node Type can declare to
1325 expose. The TOSCA Simple Profile seeks to simplify the need for declaring specific Requirement Types from
1326 nodes and instead rely upon nodes declaring their features sets using TOSCA Capability Types along with a
1327 named Feature notation.

1328 Currently, there are no use cases in this TOSCA Simple Profile in YAML specification that utilize an independently
1329 defined Requirement Type. This is a desired effect as part of the simplification of the TOSCA v1.0 specification.

1330 **A.6.8 Node Type**

1331 A Node Type is a reusable entity that defines the type of one or more Node Templates. As such, a Node Type
1332 defines the structure of observable properties via a *Properties Definition, the Requirements and Capabilities of the*
1333 *node as well as its supported interfaces.*

1334 **A.6.8.1 Keynames**

1335 The following is the list of recognized keynames for a TOSCA Node Type definition:

Keyname	Required	Definition/Type	Description
derived_from	no	string	An optional parent Node Type name this new Node Type derives from.
description	no	description	An optional description for the Node Type.
properties	no	list of property definitions	An optional list of property definitions for the Node Type.
attributes	no	list of attribute definitions	An optional list of attribute definitions for the Node Type.
requirements	no	list of requirement definitions	An optional <i>sequenced</i> list of requirement definitions for the Node Type.
capabilities	no	list of capability definitions	An optional list of capability definitions for the Node Type.
interfaces	no	list of interface definitions	An optional list of interface definitions supported by the Node Type.
artifacts	no	list of artifacts definitions	An optional list of named artifact definitions for the Node Type.

1336 **A.6.8.2 Grammar**

1337 Node Types have following grammar:

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

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

1338 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1339 • **node_type_name**: represents the required symbolic name of the Node Type being declared.
- 1340 • **parent_node_type_name**: represents the name (*string*) of the [Node Type](#) this Node Type definition
1341 derives from (i.e., its “parent” type).
- 1342 • **node_type_description**: represents the optional [description](#) string for the corresponding
1343 **node_type_name**.
- 1344 • **property_definitions**: represents the optional list of [property definitions](#) for the Node Type.
- 1345 • **attribute_definitions**: represents the optional list of [attribute definitions](#) for the Node Type.
- 1346 • **requirement_definitions**: represents the optional *sequenced* list of [requirement definitions](#) for the
1347 Node Type.
- 1348 • **capability_definitions**: represents the optional list of [capability definitions](#) for the Node Type.
- 1349 • **interface_definitions**: represents the optional list of one or more [interface definitions](#) supported by
1350 the Node Type.
- 1351 • **artifact_definitions**: represents the optional list of [artifact definitions](#) for the Node Template that
1352 augment those provided by its declared Node Type.

1353 **A.6.8.3 Additional Requirements**

- 1354 • Requirements are intentionally expressed as a sequenced list of TOSCA [Requirement definitions](#) which
1355 SHOULD be resolved (processed) in sequence order by TOSCA Orchestrators. .

1356 **A.6.8.4 Best Practices**

- 1357 • It is recommended that all Node Types SHOULD derive directly (as a parent) or indirectly (as an ancestor) of
1358 the TOSCA Root Node Type (i.e., **tosca.nodes.Root**) to promote compatibility and portability. However, it is
1359 permitted to author Node Types that do not do so.
- 1360 • TOSCA Orchestrators, having a full view of the complete application topology template and its resultant
1361 dependency graph of nodes and relationships, **MAY** prioritize how they instantiate the nodes and
1362 relationships for the application (perhaps in parallel where possible) to achieve the greatest efficiency

1363 **A.6.8.5 Example**

```
my_company.my_types.my_app_node_type:
  derived_from: tosca.nodes.SoftwareComponent
  description: My company's custom applicaton
  properties:
    my_app_password:
      type: string
      description: application password
  constraints:
```

```

    - min_length: 6
    - max_length: 10
  attributes:
    my_app_port:
      type: integer
      description: application port number
  requirements:
    - some_database:
        capability: EndPoint.Database
        node: Database
        relationship: ConnectsTo

```

1364 A.6.9 Relationship Type

1365 A Relationship Type is a reusable entity that defines the type of one or more relationships between Node Types
1366 or Node Templates.

1367 A.6.9.1 Keynames

1368 The following is the list of recognized keynames for a TOSCA Relationship Type definition:

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

1369 A.6.9.2 Grammar

1370 Relationship Types have following grammar:

```

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

```

```
valid_target_types: [ <capability_type_names> ]
```

1371 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1372 • **relationship_type_name**: represents the required symbolic name of the Relationship Type being
1373 declared as a [string](#).
- 1374 • **parent_relationship_type_name**: represents the name ([string](#)) of the [Relationship Type](#) this
1375 Relationship Type definition derives from (i.e., its “parent” type).
- 1376 • **relationship_description**: represents the optional [description](#) string for the corresponding
1377 **relationship_type_name**.
- 1378 • **property_definitions**: represents the optional list of [property definitions](#) for the Relationship Type.
- 1379 • **attribute_definitions**: represents the optional list of [attribute definitions](#) for the Relationship Type.
- 1380 • **interface_definitions**: represents the optional list of one or more names of valid [interface](#)
1381 [definitions](#) supported by the Relationship Type.
- 1382 • **capability_type_names**: represents one or more names of valid target types for the relationship (i.e.,
1383 Capability Types).

1384 **A.6.9.3 Best Practices**

- 1385 • For TOSCA application portability, it is recommended that designers use the normative Relationship
1386 types defined in this specification where possible and derive from them for customization purposes.
- 1387 • The TOSCA Root Relationship Type (**tosca.relationships.Root**) SHOULD be used to derive new types
1388 where possible when defining new relationships types. This assures that its normative configuration
1389 interface (tosca.interfaces.relationship.Configure) can be used in a deterministic way by TOSCA
1390 orchestrators.

1391 **A.6.9.4 Examples**

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

1392 **A.7 Template-specific definitions**

1393 The definitions in this section provide reusable modeling element grammars that are specific to the Node or
1394 Relationship templates.

1395 **A.7.1 Capability assignment**

1396 A capability assignment allows node template authors to assign values to properties and attributes for a named
1397 capability definition that is part of a Node Template’s type definition.

1398 **A.7.1.1 Keynames**

1399 The following is the list of recognized keynames for a TOSCA capability assignment:

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

1400 A.7.1.2 Grammar

1401 Capability assignments have one of the following grammars:

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

1402 In the above grammars, the pseudo values that appear in angle brackets have the following meaning:

- 1403 • **capability_definition_name**: represents the symbolic name of the capability as a [string](#).
- 1404 • **property_assignments**: represents the optional list of [property assignments](#) for the capability definition.
- 1405 • **attribute_assignments**: represents the optional list of [attribute assignments](#) for the capability definition.

1408 A.7.1.3 Example

1409 The following example shows a capability assignment:

1410 A.7.1.3.1 Notation example

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

1411 A.7.2 Requirement assignment

1412 A Requirement assignment allows template authors to provide either concrete names of TOSCA templates or
 1413 provide abstract selection criteria for providers to use to find matching TOSCA templates that are used to fulfill a
 1414 named requirement's declared TOSCA Node Type.

1415 A.7.2.1 Keynames

1416 The following is the list of recognized keynames for a TOSCA requirement assignment:

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

1417 The following is the list of recognized keynames for a TOSCA requirement assignment's **relationship** keyname
1418 which is used when Property assignments need to be provided to inputs of declared interfaces or their
1419 operations:

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

1420 A.7.2.2 Grammar

1421 Named requirement assignments have one of the following grammars:

1422 A.7.2.2.1 Short notation:

1423 The following single-line grammar may be used if only a concrete Node Template for the target node needs to
1424 be declared in the requirement:

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

1425 This notation is only valid if the corresponding Requirement definition in the Node Template's parent Node Type
1426 declares (at a minimum) a valid Capability Type which can be found in the declared target Node Template. A
1427 valid capability definition always needs to be provided in the requirement declaration of the *source* node to
1428 identify a specific capability definition in the *target* node the requirement will form a TOSCA relationship with.

1429 A.7.2.2.2 Extended notation:

1430 The following grammar would be used if the requirement assignment needs to provide more information than
1431 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 ]

```

1432 **A.7.2.2.3 Extended grammar with Property Assignments for the relationship's Interfaces**

1433 The following additional multi-line grammar is provided for the relationship keyname in order to provide new
 1434 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>

```

1435 Examples of uses for the extended requirement assignment grammar include:

- 1436 • The need to allow runtime selection of the target node based upon an abstract Node Type rather than a
 1437 concrete Node Template. This may include use of the node_filter keyname to provide node and
 1438 capability filtering information to find the “best match” of a concrete Node Template at runtime.
- 1439 • The need to further clarify the concrete Relationship Template or abstract Relationship Type to use
 1440 when relating the source node’s requirement to the target node’s capability.
- 1441 • The need to further clarify the concrete capability (symbolic) name or abstract Capability Type in the
 1442 target node to form a relationship between.
- 1443 • The need to (further) constrain the occurrences of the requirement in the instance model.

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

- 1445 • **requirement_name**: represents the symbolic name of a requirement assignment as a [string](#).
- 1446 • **node_template_name**: represents the optional name of a Node Template that contains the capability
 1447 this requirement will be fulfilled by.
- 1448 • **relationship_template_name**: represents the optional name of a [Relationship Type](#) to be used when
 1449 relating the requirement appears to the capability in the target node.
- 1450 • **capability_symbolic_name**: represents the optional ordered list of specific, required capability type or
 1451 named capability definition within the target Node Type or Template.
- 1452 • **node_type_name**: represents the optional name of a TOSCA Node Type the associated named
 1453 requirement can be fulfilled by. This must be a type that is compatible with the Node Type declared on
 1454 the matching requirement (same symbolic name) the requirement’s Node Template is based upon.

- 1455 • **relationship_type_name**: represents the optional name of a [Relationship Type](#) that is compatible with
1456 the Capability Type in the target node.
- 1457 • **property_assignments**: represents the optional list of property value assignments for the declared
1458 relationship.
- 1459 • **interface_assignments**: represents the optional list of interface definitions for the declared
1460 relationship used to provide property assignments on inputs of interfaces and operations.
- 1461 • **capability_type_name**: represents the optional name of a Capability Type definition within the target
1462 Node Type this requirement needs to form a relationship with.
- 1463 • **node_filter_definition**: represents the optional [node filter](#) TOSCA orchestrators would use to fulfill
1464 the requirement for selecting a target node. Note that this SHALL only be valid if the **node** keyname's
1465 value is a Node Type and is invalid if it is a Node Template.

1466 **A.7.2.3 Examples**

1467 **A.7.2.3.1 Example 1 – Abstract hosting requirement on a Node Type**

1468 A web application node template named 'my_application_node_template' of type **WebApplication** declares
1469 a requirement named 'host' that needs to be fulfilled by any node that derives from the node type **WebServer**.

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

1470 In this case, the node template's type is **WebApplication** which already declares the Relationship Type
1471 **HostedOn** to use to relate to the target node and the Capability Type of **Container** to be the specific target of
1472 the requirement in the target node.

1473 **A.7.2.3.2 Example 2 - Requirement with Node Template and a custom Relationship Type**

1474 This example is similar to the previous example; however, the requirement named 'database' describes a
1475 requirement for a connection to a database endpoint (**Endpoint.Database**) Capability Type in a named node
1476 template (**my_database**). However, the connection requires a custom Relationship Type
1477 (**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
```

1478 **A.7.2.3.3 Example 3 - Requirement for a Compute node with additional selection criteria (filter)**

1479 This example shows how to extend an abstract 'host' requirement for a **Compute** node with a filter definition that
1480 further constrains TOSCA orchestrators to include additional properties and capabilities on the target node when
1481 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: 2 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 ] }
```

1482 **A.7.3 Node Template**

1483 A Node Template specifies the occurrence of a manageable software component as part of an application's
1484 topology model which is defined in a TOSCA Service Template. A Node template is an instance of a specified
1485 Node Type and can provide customized properties, constraints or operations which override the defaults provided
1486 by its Node Type and its implementations.

1487 **A.7.3.1 Keynames**

1488 The following is the list of recognized keynames for a TOSCA Node Template definition:

Keyname	Required	Type	Description
type	yes	string	The required name of the Node Type the Node Template is based upon.
description	no	description	An optional description for the Node Template.
directives	no	string[]	An optional list of directive values to provide processing instructions to orchestrators and tooling.
properties	no	list of property assignments	An optional list of property value assignments for the Node Template.
attributes	no	list of attribute assignments	An optional list of attribute value assignments for the Node Template.

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

1489

A.7.3.2 Grammar

```

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

```

1490

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

1491

- **node_template_name**: represents the required symbolic name of the Node Template being declared.

1492

- **node_type_name**: represents the name of the Node Type the Node Template is based upon.

1493

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

1494

- **directives**: represents the optional list of processing instruction keywords for treatment of the node

1495

for tooling and orchestrators.

- 1496 • **property_assignments**: represents the optional list of [property assignments](#) for the Node Template
1497 that provide values for properties defined in its declared Node Type.
- 1498 • **attribute_assignments**: represents the optional list of [attribute assignments](#) for the Node Template
1499 that provide values for attributes defined in its declared Node Type.
- 1500 • **requirement_assignments**: represents the optional *sequenced* list of [requirement assignments](#) for the
1501 Node Template that allow assignment of type-compatible capabilities, target nodes, relationships and
1502 target (node filters) for use when fulfilling the requirement at runtime.
- 1503 • **capability_assignments**: represents the optional list of [capability assignments](#) for the Node Template
1504 that augment those provided by its declared Node Type.
- 1505 • **interface_definitions**: represents the optional list of [interface definitions](#) for the Node Template
1506 that augment those provided by its declared Node Type.
- 1507 • **artifact_definitions**: represents the optional list of [artifact definitions](#) for the Node Template that
1508 augment those provided by its declared Node Type.
- 1509 • **node_filter_definition**: represents the optional [node filter](#) TOSCA orchestrators would use for
1510 selecting a matching node template.
- 1511 • **source_node_template_name**: represents the optional (symbolic) name of another node template to
1512 copy into (all keynames and values) and use as a basis for this node template.

1513 A.7.3.3 Additional requirements

- 1514 • The **node_filter** keyword (and supporting grammar) **SHALL** only be valid if the Node Template has a
1515 **directive** keyname with the value of “**selectable**” set.
- 1516 • The source node template provided as a value on the **copy** keyname **MUST NOT** itself use the **copy**
1517 keyname (i.e., it must itself be a complete node template description and not copied from another node
1518 template).

1519 A.7.3.4 Example

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

1520 A.7.4 Relationship Template

1521 A Relationship Template specifies the occurrence of a manageable relationship between node templates as part
1522 of an application’s topology model which is defined in a TOSCA Service Template. A Relationship template is an
1523 instance of a specified Relationship Type and can provide customized properties, constraints or operations which
1524 override the defaults provided by its Relationship Type and its implementations.

1525 The following is the list of recognized keynames for a TOSCA Relationship Template definition:

Keyname	Required	Type	Description
type	yes	string	The required name of the Relationship Type the Relationship Template is based upon.
alias	no	string	The optional name of a different Relationship Template definition whose values are (effectively) copied into the definition for this Relationship Template (prior to any other overrides).
description	no	description	An optional description for the Relationship Template.
properties	no	list of property assignments	An optional list of property assignments for the Relationship Template.
attributes	no	list of attribute assignments	An optional list of attribute assignments for the Relationship Template.
interfaces	no	list of interface definitions	An optional list of named interface definitions for the Node Template.
copy	no	name	The optional (symbolic) name of another relationship template to copy into (all keynames and values) and use as a basis for this relationship template.

1526 A.7.4.1 Grammar

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

1527 In the above grammar, the pseudo values that appear in angle brackets have the following meaning:

- 1528 • **relationship_template_name**: represents the required symbolic name of the Relationship Template
1529 being declared.
- 1530 • **relationship_type_name**: represents the name of the Relationship Type the Relationship Template is
1531 based upon.
- 1532 • **relationship_template_description**: represents the optional **description** string for the Relationship
1533 Template.
- 1534 • **property_assignments**: represents the optional list of **property assignments** for the Relationship
1535 Template that provide values for properties defined in its declared Relationship Type.
- 1536 • **attribute_assignments**: represents the optional list of **attribute assignments** for the Relationship
1537 Template that provide values for attributes defined in its declared Relationship Type.
- 1538 • **interface_definitions**: represents the optional list of **interface definitions** for the Relationship
1539 Template that augment those provided by its declared Relationship Type.

- 1540
- **source_relationship_template_name**: represents the optional (symbolic) name of another relationship template to copy into (all keynames and values) and use as a basis for this relationship template.
- 1541
- 1542

1543 **A.7.4.2 Additional requirements**

- The source relationship template provided as a value on the **copy** keyname MUST NOT itself use the **copy** keyname (i.e., it must itself be a complete relationship template description and not copied from another relationship template).
- 1544
- 1545
- 1546

1547 **A.7.4.3 Example**

```
relationship_templates:
```

1548 **A.8 Topology Template definition**

1549 This section defines the topology template of a cloud application. The main ingredients of the topology template
1550 are node templates representing components of the application and relationship templates representing links
1551 between the components. These elements are defined in the nested **node_templates** section and the nested
1552 **relationship_templates** sections, respectively. Furthermore, a topology template allows for defining input
1553 parameters, output parameters as well as grouping of node templates.

1554 **A.8.1 Grammar**

1555 The overall grammar of the **topology_template** section is shown below.–Detailed grammar definitions of the
1556 each sub-sections are provided in subsequent subsections.

```
topology_template:  
  description:  
    # a description of the topology template  
  
  inputs:  
    # definition of input parameters for the topology template  
  
  node_templates:  
    # definition of the node templates of the topology  
  
  relationship_templates:  
    # definition of the relationship templates of the topology  
  
  outputs:  
    # definition of output parameters for the topology template  
  
  groups:  
    # definition of logical groups of node templates within the topology  
  
  substitution_mappings:
```

```
node_type: <node_type_name>
capabilities:
  <map_of_capability_mappings_to_expose>
requirements:
  <map_of_requirement_mapping_to_expose>
```

1557 **A.8.1.1 inputs**

1558 The **inputs** section provides a means to define parameters using TOSCA property definitions, their allowed
1559 values via constraints and default values within a TOSCA Simple Profile template. Input parameters defined in the
1560 **inputs** section of a topology template can be mapped to properties of node templates or relationship templates
1561 within the same topology template and can thus be used for parameterizing the instantiation of the topology
1562 template.

1563
1564 This section defines topology template-level input parameter section.

- 1565 • Inputs here would ideally be mapped to BoundaryDefinitions in TOSCA v1.0.
- 1566 • Treat input parameters as fixed global variables (not settable within template)
- 1567 • If not in input take default (nodes use default)

1568 **A.8.1.1.1 Grammar**

1569 The grammar of the **inputs** section is as follows:

```
inputs:
  <property_definition_1>
  ...
  <property_definition_n>
```

1570 **A.8.1.1.2 Examples**

1571 This section provides a set of examples for the single elements of a topology template.

1572 Simple **inputs** example without any constraints:

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

1573 Example of **inputs** with constraints:

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

```
- min_length: 9
```

1574 **A.8.1.2 node_templates**

1575 The **node_templates** section lists the Node Templates that describe the (software) components that are used to
1576 compose cloud applications.

1577 **A.8.1.2.1 grammar**

1578 The grammar of the **node_templates** section is as follows:

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

1579 **A.8.1.2.2 Example**

1580 Example of **node_templates** section:

```
node_templates:  
  my_webapp_node_template:  
    type: WebApplication  
  
  my_database_node_template:  
    type: Database
```

1581 **A.8.1.3 relationship_templates**

1582 The **relationship_templates** section lists the Relationship Templates that describe the relations between
1583 components that are used to compose cloud applications.

1584

1585 Note that in the TOSCA Simple Profile, the explicit definition of relationship templates as it was required in
1586 TOSCA v1.0 is optional, since relationships between nodes get implicitly defined by referencing other node
1587 templates in the requirements sections of node templates.

1588 **A.8.1.3.1 Grammar**

1589 The grammar of the **relationship_templates** section is as follows:

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

1590 **A.8.1.3.2 Example**

1591 Example of **relationship_templates** section:

```
relationship_templates:  
  my_connectsto_relationship:
```

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

1592 **A.8.1.4 outputs**

1593 The **outputs** section provides a means to define the output parameters that are available from a TOSCA Simple
1594 Profile service template. It allows for exposing attributes of node templates or relationship templates within the
1595 containing **topology_template** to users of a service.

1596 **A.8.1.4.1 Grammar**

1597 The grammar of the **outputs** section is as follows:

```
outputs:
  <attribute_assignments>
```

1598 **A.8.1.4.2 Example**

1599 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 ] }
```

1600 **A.8.1.5 groups**

1601 The **groups** section allows for grouping one or more node templates within a TOSCA Service Template and for
1602 assigning special attributes like policies to the group.

1603 **A.8.1.5.1 Grammar**

1604 The grammar of the **groups** section is as follows:

```
groups:
  <group_symbolic_name_1>:
    members: [ node_template_name_1, ..., node_template_name_n ]
    policies:
      - <optional_list of policy_names_for_group_1>
    ...
  <group_symbolic_name_N>
    members: [ node_template_name_1, ..., node_template_name_n ]
    policies:
      - <optional_list of policy_names_for_group_N>
```

1605 **A.8.1.5.2 Example**

1606 The following example shows the definition of three Compute nodes in the **node_templates** section of a
1607 **topology_template** as well as the grouping of two of the Compute nodes in a group **server_group_1**.

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

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

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

groups:
  server_group_1:
    members: [ server1, server2 ]
    policies:
      - anti_collocation_policy:
          # specific policy declarations omitted, as this is not yet specified
```

1608 **A.8.2 Notes**

- 1609
- 1610 • The parameters (properties) that are listed as part of the **inputs** block can be mapped to
1611 **PropertyMappings** provided as part of **BoundaryDefinitions** as described by the TOSCA v1.0
specification.
 - 1612 • The node templates listed as part of the **node_templates** block can be mapped to the list of
1613 **NodeTemplate** definitions provided as part of **TopologyTemplate** of a **ServiceTemplate** as described
1614 by the TOSCA v1.0 specification.
 - 1615 • The relationship templates listed as part of the **relationship_templates** block can be mapped to the
1616 list of **RelationshipTemplate** definitions provided as part of **TopologyTemplate** of a
1617 **ServiceTemplate** as described by the TOSCA v1.0 specification.
 - 1618 • The output parameters that are listed as part of the **outputs** section of a topology template can be
1619 mapped to **PropertyMappings** provided as part of **BoundaryDefinitions** as described by the TOSCA
1620 v1.0 specification.
 - 1621 ○ Note, however, that TOSCA v1.0 does not define a direction (input vs. output) for those
1622 mappings, i.e. TOSCA v1.0 **PropertyMappings** are underspecified in that respect and TOSCA
1623 Simple Profile's **inputs** and **outputs** provide a more concrete definition of input and output
1624 parameters.

1625 A.9 Service Template definition

1626 A TOSCA Service Template (YAML) document contains element definitions of building blocks for cloud
1627 application, or complete models of cloud applications. This section describes the top-level structural elements
1628 (TOSCA keynames) along with their grammars, which are allowed to appear in a TOSCA Service Template
1629 document.

1630 A.9.1 Keynames

1631 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.
tosca_default_namespace	no	string	Defines the namespace of the TOSCA schema to use for validation.
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.
imports	no	list of string	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.
dsl_defintions	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.
data_types	no	list of Data Types	Declares a list of optional TOSCA Data Type definitions.
node_types	no	list of Node Types	This section contains a set of node type definitions for use in service templates.
relationship_types	no	list of Relationship Types	This section contains a set of relationship type definitions for use in service templates.
capability_types	no	list of Capability Types	This section contains an optional list of capability type definitions for use in service templates.
artifact_types	no	list of Artifact Types	This section contains an optional list of artifact type definitions for use in service templates
interface_types	no	list of Interface Types	This section contains an optional list of interface type definitions for use in service templates.
topology_template	no	Topology Template	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.

1632 A.9.1.1 Metadata keynames

1633 The following is the list of recognized metadata keynames for a TOSCA Service Template definition:

Keyname	Required	Type	Description
template_name	no	string	Declares a descriptive name for the template.
template_author	no	string	Declares the author(s) or owner of the template.
template_version	no	string	Declares the version string for the template.

1634 **A.9.2 Grammar**

1635 The overall structure of a TOSCA Service Template and its top-level key collations using the TOSCA Simple
1636 Profile is shown below:

```

tosca_definitions_version: # Required TOSCA Definitions version string
tosca_default_namespace:  # Optional. default namespace (for type schema)

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

# Optional description of the definitions inside the file.
description: <template_type_description>

imports:
  # list of import statements for importing other definitions files

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

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

data_types:
  # list of TOSCA datatype definitions

node_types:
  # list of node type definitions

capability_types:
  # list of capability type definitions

relationship_types:
  # list of relationship type definitions

```

```
artifact_types:
  # list of artifact type definitions

interface_types
  # list of interface type definitions

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

1637 **A.9.2.1 Notes**

- 1638 • TOSCA Service Templates do not have to contain a topology_template and MAY contain simply type
1639 definitions (e.g., Artifact, Interface, Capability, Node, Relationship Types, etc.) and be imported for use
1640 as type definitions in other TOSCA Service Templates.

1641 **A.9.3 Top-level keyname definitions**

1642 **A.9.3.1 tosca_definitions_version**

1643 This required element provides a means to include a reference to the TOSCA Simple Profile specification within
1644 the TOSCA Definitions YAML file. It is an indicator for the version of the TOSCA grammar that should be used to
1645 parse the remainder of the document.

1646 **A.9.3.1.1 Keyname**

```
tosca_definitions_version
```

1647 **A.9.3.1.2 Grammar**

1648 Single-line form:

```
tosca_definitions_version: <tosca_simple_profile_version>
```

1649 **A.9.3.1.3 Examples:**

1650 TOSCA Simple Profile version 1.0 specification using the defined namespace alias (see Section A.1):

```
tosca_definitions_version: tosca_simple_yaml_1_0_0
```

1651 TOSCA Simple Profile version 1.0 specification using the fully defined (target) namespace (see Section A.1):

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

1652 **A.9.3.2 template_name**

1653 This optional element declares the optional name of service template as a single-line string value.

1654 **A.9.3.2.1 Keyname**

```
template_name
```

1655 **A.9.3.2.2 Grammar**

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

1656 **A.9.3.2.3 Example**

```
template_name: My service template
```

1657 **A.9.3.2.4 Notes**

- 1658 • Some service templates are designed to be referenced and reused by other service templates.
1659 Therefore, in these cases, the **template_name** value SHOULD be designed to be used as a unique
1660 identifier through the use of namespacing techniques.

1661 **A.9.3.3 template_author**

1662 This optional element declares the optional author(s) of the service template as a single-line string value.

1663 **A.9.3.3.1 Keyname**

```
template_author
```

1664 **A.9.3.3.2 Grammar**

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

1665 **A.9.3.3.3 Example**

```
template_author: My service template
```

1666 **A.9.3.4 template_version**

1667 This element declares the optional version of the service template as a single-line string value.

1668 **A.9.3.4.1 Keyname**

```
template_version
```

1669 **A.9.3.4.2 Grammar**

```
template_version: <version>
```

1670 **A.9.3.4.3 Example**

```
template_version: 2.0.17
```

1671 **A.9.3.4.4 Notes:**

- 1672 • Some service templates are designed to be referenced and reused by other service templates and have
1673 a lifecycle of their own. Therefore, in these cases, a **template_version** value SHOULD be included and
1674 used in conjunction with a unique **template_name** value to enable lifecycle management of the service
1675 template and its contents.

1676 **A.9.3.5 description**

1677 This optional element provides a means to include single or multiline descriptions within a TOSCA Simple Profile
1678 template as a scalar string value.

1679 **A.9.3.5.1 Keyname**

```
description
```

1680 **A.9.3.6 imports**

1681 This optional element provides a way to import a *block sequence* of one or more TOSCA Definitions documents.
1682 TOSCA Definitions documents can contain reusable TOSCA type definitions (e.g., Node Types, Relationship
1683 Types, Artifact Types, etc.) defined by other authors. This mechanism provides an effective way for companies
1684 and organizations to define normative types and/or describe their software applications for reuse in other TOSCA
1685 Service Templates.

1686 **A.9.3.6.1 Keyname**

```
imports
```

1687 **A.9.3.6.2 Grammar**

```
imports:  
  - <tosca_definitions_file_URI_1>  
  - ...  
  - <tosca_definitions_file_URI_n>
```

1688 **A.9.3.6.3 Example**

```
# An example import of definitions files from a location relative to the  
# file location of the service template declaring the import.  
imports:  
  - relative_path/my_defns/my_typesdefs_1.yaml  
  - ...  
  - relative_path/my_defns/my_typesdefs_n.yaml
```

1689 **A.9.3.7 dsl_definitions**

1690 This optional element provides a section to define macros (e.g., YAML-style macros when using the TOSCA
1691 Simple Profile in YAML specification).

1692 **A.9.3.7.1 Keyname**

```
dsl_definitions
```

1693 **A.9.3.7.2 Grammar**

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

1694 **A.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
```

1695 **A.9.3.8 datatype_definitions**

1696 This optional element provides a section to define new datatypes in TOSCA.

1697 **A.9.3.8.1 Keyname**

```
datatype_definitions
```

1698 **A.9.3.8.2 Grammar**

```
datatype_definitions:
  <tosca_datatype_def_1>
  ...
  <tosca_datatype_def_n>
```

1699 **A.9.3.8.3 Example**

```
datatype_definitions:
  # 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
```

1700 **A.9.3.9 node_types**

1701 This element lists the Node Types that provide the reusable type definitions for software components that Node
1702 Templates can be based upon.

1703 **A.9.3.9.1 Keyname**

```
node_types
```

1704 **A.9.3.9.2 Grammar**

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

1705 **A.9.3.9.3 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
```

1706 **A.9.3.9.4 Notes**

- 1707
- The node types listed as part of the **node_types** block can be mapped to the list of **NodeType** definitions
1708 as described by the TOSCA v1.0 specification.

1709 **A.9.3.10 relationship_types**

1710 This element lists the Relationship Types that provide the reusable type definitions that can be used to describe
1711 dependent relationships between Node Templates or Node Types.

1712 **A.9.3.10.1 Keyname**

```
relationship_types
```

1713 **A.9.3.10.2 Grammar**

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

1714 **A.9.3.10.3 Example**

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

1715 **A.9.3.11 capability_types**

1716 This element lists the Capability Types that provide the reusable type definitions that can be used to describe
1717 features Node Templates or Node Types can declare they support.

1718 **A.9.3.11.1 Keyname**

```
capability_types
```

1719 **A.9.3.11.2 Grammar**

```
capability_types:  
  <capability_type_defn_1>  
  ...  
  <capability type_defn_n>
```

1720 **A.9.3.11.3 Example**

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

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

Appendix B. Functions

This section includes functions that are supported for use within a TOSCA Service Template.

B.1 Reserved Function Keywords

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

Keyword	Valid Contexts	Description
SELF	Node Template or Relationship Template	A TOSCA orchestrator will interpret this keyword as the Node or Relationship Template instance that contains the function at the time the function is evaluated.
SOURCE	Relationship Template only.	A TOSCA orchestrator will interpret this keyword as the Node Template instance that is at the source end of the relationship that contains the referencing function.
TARGET	Relationship Template only.	A TOSCA orchestrator will interpret this keyword as the Node Template instance that is at the target end of the relationship that contains the referencing function.
HOST	Node Template only	<p>A TOSCA orchestrator will interpret this keyword to refer to the all nodes that “host” the node using this reference (i.e., as identified by its HostedOn relationship).</p> <p>Specifically, TOSCA orchestrators that encounter this keyword when evaluating the get_attribute or get_property functions SHALL search each node along the “HostedOn” relationship chain starting at the immediate node that hosts the node where the function was evaluated (and then that node’s host node, and so forth) until a match is found or the “HostedOn” relationship chain ends.</p>

B.2 Environment Variable Conventions

B.2.1 Reserved Environment Variable Names and Usage

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

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

Keyword	Valid Contexts	Description
TARGETS	Relationship Template only.	<ul style="list-style-type: none">For an implementation artifact that is executed in the context of a relationship, this keyword, if present, is used to supply a list of Node Template instances in a TOSCA application’s instance model that are currently target of the context relationship.The value of this environment variable will be a comma-separated list of identifiers of the single target node instances (i.e., the tosca_id attribute of the node).

Keyword	Valid Contexts	Description
TARGET	Relationship Template only.	<ul style="list-style-type: none"> For an implementation artifact that is executed in the context of a relationship, this keyword, if present, identifies a Node Template instance in a TOSCA application's instance model that is a target of the context relationship, and which is being acted upon in the current operation. The value of this environment variable will be the identifier of the single target node instance (i.e., the <code>tosca_id</code> attribute of the node).
SOURCES	Relationship Template only.	<ul style="list-style-type: none"> For an implementation artifact that is executed in the context of a relationship, this keyword, if present, is used to supply a list of Node Template instances in a TOSCA application's instance model that are currently source of the context relationship. The value of this environment variable will be a comma-separated list of identifiers of the single source node instances (i.e., the <code>tosca_id</code> attribute of the node).
SOURCE	Relationship Template only.	<ul style="list-style-type: none"> For an implementation artifact that is executed in the context of a relationship, this keyword, if present, identifies a Node Template instance in a TOSCA application's instance model that is a source of the context relationship, and which is being acted upon in the current operation. The value of this environment variable will be the identifier of the single source node instance (i.e., the <code>tosca_id</code> attribute of the node).

1741

1742 For scripts (or implementation artifacts in general) that run in the context of relationship operations, select
 1743 properties and attributes of both the relationship itself as well as select properties and attributes of the source and
 1744 target node(s) of the relationship can be provided to the environment by declaring respective operation inputs.
 1745

1746

1746 Declared inputs from mapped properties or attributes of the source or target node (selected via the **SOURCE** or
 1747 **TARGET** keyword) will be provided to the environment as variables having the exact same name as the inputs. In
 1748 addition, the same values will be provided for the complete set of source or target nodes, however prefixed with
 1749 the ID if the respective nodes. By means of the **SOURCES** or **TARGETS** variables holding the complete set of source
 1750 or target node IDs, scripts will be able to iterate over corresponding inputs for each provided ID prefix.
 1751

1752

1752 The following example snippet shows an imaginary relationship definition from a load-balancer node to worker
 1753 nodes. A script is defined for the **add_target** operation of the Configure interface of the relationship, and the
 1754 **ip_address** attribute of the target is specified as input to the script:
 1755

1755

```

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
  
```

1756 The **add_target** operation will be invoked, whenever a new target member is being added to the load-balancer.
1757 With the above inputs declaration, a **member_ip** environment variable that will hold the IP address of the target
1758 being added will be provided to the **configure_members.py** script. In addition, the IP addresses of all current
1759 load-balancer members will be provided as environment variables with a naming scheme of **<target node**
1760 **ID>_member_ip**. This will allow, for example, scripts that always just write the complete list of load-balancer
1761 members into a configuration file to do so instead of updating existing list, which might be more complicated.

1762 Assuming that the TOSCA application instance includes five load-balancer members, **node1** through **node5**,
1763 where **node5** is the current target being added, the following environment variables (plus potentially more
1764 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
```

1765 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 ...
```

1766 **B.2.2 Prefixed vs. Unprefixed TARGET names**

1767 The list target node types assigned to the TARGETS key in an execution environment would have names prefixed
1768 by unique IDs that distinguish different instances of a node in a running model. Future drafts of this specification
1769 will show examples of how these names/IDs will be expressed.

1770 **B.2.2.1 Notes**

- 1771 • Target of interest is always un-prefixed. Prefix is the target opaque ID. The IDs can be used to find the
1772 environment var. for the corresponding target. Need an example here.
- 1773 • If you have one node that contains multiple targets this would also be used (add or remove target
1774 operations would also use this you would get set of all current targets).

1775 **B.3 Intrinsic functions**

1776 These functions are supported within the TOSCA template for manipulation of template data.

1777 **B.3.1 concat**

1778 The **concat** function is used to concatenate two or more string values within a TOSCA service template.

1779 **B.3.1.1 Grammar**

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

1780 **B.3.1.2 Parameters**

Parameter	Required	Type	Description
<string_value_expressions_*>	yes	list of string or string value expressions	A list of one or more strings (or expressions that result in a string value) which can be concatenated together into a single string.

1781 **B.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 ] ] }
```

1782 **B.3.1 token**

1783 The **token** function is used within a TOSCA service template on a string to parse out (tokenize) substrings separated by one or more token characters within a larger string.

1785 **B.3.1.1 Grammar**

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

1786 **B.3.1.2 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.

1787 **B.3.1.3 Examples**

```
outputs:
  webserver_port:
    description: the port provided at the end of my server's endpoint's IP address
    value: { token: [ get_attribute: [ my_server, data_endpoint, ip_address ],
                  ':',
                  1 ] }
```

1788 **B.4 Property functions**

1789 These functions are used within a service template to obtain property values from property definitions declared
1790 elsewhere in the same service template. These property definitions can appear either directly in the service
1791 template itself (e.g., in the inputs section) or on entities (e.g., node or relationship templates) that have been
1792 modeled within the template.

1793

1794 Note that the **get_input** and **get_property** functions may only retrieve the static values of property definitions of
1795 a TOSCA application as defined in the TOSCA Service Template. The **get_attribute** function should be used
1796 to retrieve values for attribute definitions (or property definitions reflected as attribute definitions) from the runtime
1797 instance model of the TOSCA application (as realized by the TOSCA orchestrator).

1798 **B.4.1 get_input**

1799 The **get_input** function is used to retrieve the values of properties declared within the **inputs** section of a
1800 TOSCA Service Template.

1801 **B.4.1.1 Grammar**

```
get_input: <input_property_name>
```

1802 **B.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.

1803 **B.4.1.3 Examples**

```
inputs:
  cpus:
    type: integer

node_templates:
  my_server:
    type: toska.nodes.Compute
    capabilities:
      host:
        properties:
```

```
num_cpus: { get_input: cpus }
```

1804 B.4.2 get_property

1805 The `get_property` function is used to retrieve property values between modelable entities defined in the same
1806 service template.

1807 B.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> ]
```

1808 B.4.2.2 Parameters

Parameter	Required	Type	Description
<modelable_entity_name> SELF SOURCE TARGET HOST	yes	string	The required name of a modelable entity (e.g., Node Template or Relationship Template name) as declared in the service template that contains the named property definition the function will return the value from. See section B.1 for valid keywords.
<optional_req_or_cap_name>	no	string	The optional name of the requirement or capability name within the modelable entity (i.e., the <modelable_entity_name> which contains the named property definition the function will return the value from. Note: If the property definition is located in the modelable entity directly, then this parameter MAY be omitted.
<property_name>	yes	string	The name of the property definition the function will return the value from.
<nested_property_name_or_index_*>	no	string integer	Some TOSCA properties are complex (i.e., composed as nested structures). These parameters are used to dereference into the names of these nested structures when needed. Some properties represent list types. In these cases, an index may be provided to reference a specific entry in the list (as named in the previous parameter) to return.

1809 B.4.2.3 Examples

1810 The following example shows how to use the `get_property` function with an actual Node Template name:

```
node_templates:  
  
  mysql_database:  
    type: toscanodes.Database  
    properties:  
      name: sql_database1  
  
  wordpress:  
    type: toscanodes.WebApplication.WordPress
```

```

...
interfaces:
  Standard:
    configure:
      inputs:
        wp_db_name: { get_property: [ mysql_database, name ] }

```

1811 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 ] }

```

1812 The following example shows how to use the get_property function using the TARGET keyword:

TBD

1813 B.5 Attribute functions

1814 These functions (attribute functions) are used within an instance model to obtain attribute values from instances of
 1815 nodes and relationships that have been created from an application model described in a service template. The
 1816 instances of nodes or relationships can be referenced by their name as assigned in the service template or
 1817 relative to the context where they are being invoked.

1818 **B.5.1 get_attribute**

1819 The `get_attribute` function is used to retrieve the values of named attributes declared by the referenced node
1820 or relationship template name.

1821

1822 **B.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>, ]
```

1823 **B.5.1.2 Parameters**

Parameter	Required	Type	Description
<modelable_entity_name> SELF SOURCE TARGET HOST	yes	string	The required name of a modelable entity (e.g., Node Template or Relationship Template name) as declared in the service template that contains the named attribute definition the function will return the value from. See section B.1 for valid keywords.
<optional_req_or_cap_name>	no	string	The optional name of the requirement or capability name within the modelable entity (i.e., the <modelable_entity_name> which contains the named attribute definition the function will return the value from. Note: If the attribute definition is located in the modelable entity directly, then this parameter MAY be omitted.
<attribute_name>	yes	string	The name of the attribute definition the function will return the value from.
<nested_attribute_name_or_index_*>	no	string integer	Some TOSCA attributes are complex (i.e., composed as nested structures). These parameters are used to dereference into the names of these nested structures when needed. Some attributes represent list types. In these cases, an index may be provided to reference a specific entry in the list (as named in the previous parameter) to return.

1824 **B.5.1.3 Examples:**

1825 The attribute functions are used in the same way as the equivalent Property functions described above. Please
1826 see their examples and replace “get_property” with “get_attribute” function name.

1827 **B.5.1.4 Notes**

1828 These functions are used to obtain attributes from instances of node or relationship templates by the names they
1829 were given within the service template that described the application model (pattern).

1830 **B.5.1.4.1 Notes:**

- 1831 • These functions only work when the orchestrator can resolve to a single node or relationship instance for
1832 the named node or relationship. This essentially means this is acknowledged to work only when the node
1833 or relationship template being referenced from the service template has a cardinality of 1 (i.e., there can
1834 only be one instance of it running).

1835 B.6 Operation functions

1836 These functions are used within an instance model to obtain values from interface operations. These can be used
1837 in order to set an attribute of a node instance at runtime or to pass values from one operation to another.

1838 B.6.1 get_operation_output

1839 The `get_operation_output` function is used to retrieve the values of variables exposed / exported from an
1840 interface operation.

1841 B.6.1.1 Grammar

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

1842 B.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.

1843 B.6.1.3 Notes

- 1844
- If operation failed, then ignore its outputs. Orchestrators should allow orchestrators to continue
1845 running when possible past deployment in the lifecycle. For example, if an update fails, the application
1846 should be allowed to continue running and some other method would be used to alert administrators of
1847 the failure.

1848 B.7 Navigation functions

- 1849
- This version of the TOSCA Simple Profile does not define any model navigation functions.

1850 B.7.1 get_nodes_of_type

1851 The `get_nodes_of_type` function can be used to retrieve a list of all known instances of nodes of the declared
1852 Node Type.

1853 B.7.1.1 Grammar

```
get_nodes_of_type: <node_type_name>
```

1854 **B.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.

1855 **B.7.1.3 Returns**

1856 **B.7.1.4**

Return Key	Type	Description
TARGETS	<see above>	The list of node instances from the current application instance that match the node_type_name supplied as an input parameter of this function.

1857 **B.8 Artifact functions**

1858 **B.8.1 get_artifact**

1859 The **get_artifact** function is used to retrieve artifact location between modelable entities defined in the same
 1860 service template.

1861 **B.8.1.1 Grammar**

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

1862 **B.8.1.2 Parameters**

Parameter	Required	Type	Description
<modelable entity name> SELF SOURCE TARGET HOST	yes	string	The required name of a modelable entity (e.g., Node Template or Relationship Template name) as declared in the service template that contains the named property definition the function will return the value from. See section B.1 for valid keywords.
<artifact_name>	yes	string	The name of the artifact definition the function will return the value from.
<location> LOCAL_FILE	no	string	Location value must be either a valid path e.g. '/etc/var/my_file' or ' LOCAL_FILE '. If the value is LOCAL_FILE the orchestrator is responsible for providing a path as the result of the get_artifact call where the artifact file can be accessed. The orchestrator will also remove the artifact from this location at the end of the operation. If the location is a path specified by the user the orchestrator is responsible to copy the artifact to the specified location. The orchestrator will return the path as the value of the get_artifact function and leave the file here after the execution of the operation.

Parameter	Required	Type	Description
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>

1863 **B.8.1.3 Examples**

1864 The following example shows how to use the `get_artifact` function with an actual Node Template name:

1865 **B.8.1.3.1 Example: Retrieving artifact without specified location:**

```
node_templates:

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

1866 In such implementation the tosca orchestrator may provide the wordpress.zip archive as a local url (example:
 1867 <file:///home/user/wordpress.zip>) or a remote one (example: <http://cloudrepo:80/files/wordpress.zip>) (some
 1868 orchestrator may indeed provide some global artifact repository management features)

1869 **B.8.1.3.2 Example: Retrieving artifact as a local path :**

1870 The following example explains how to force the orchestrator to copy the fille locally before calling the operation's
 1871 implementation script:

1872

```
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

```

1873 In such implementation the tosca orchestrator must provide the wordpress.zip archive as a local path (example:
 1874 [/tmp/wordpress.zip](#)) and **will remove it** after the operation is completed.

1875 ***B.8.1.3.3 Example: Retrieving artifact in a specified location:***

1876 The following example explains how to force the orchestrator to copy the file locally to a specific location before
 1877 calling the operation's implementation script :

1878

```

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

```

1879 In such implementation the tosca orchestrator must provide the wordpress.zip archive as a local path (example:
 1880 [C:/wpdata/wp.zip](#)) and **will let it** after the operation is completed.

1881 **B.9 Context-based Entity name (global)**

1882 TBD

1883

1884 Goal:

- 1885 • Using the full paths of modelable entity names to qualify context with the future goal of a more robust
 1886 get_attribute function: e.g., `get_attribute(<context-based-entity-name>, <attribute name>)`

Appendix C. TOSCA normative type definitions

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.

C.1 Assumptions

- 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.

C.2 Data Types

C.2.1 `tosca.datatypes.Credential`

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

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

C.2.1.1 Properties

Name	Required	Type	Constraints	Description
protocol	yes	string	None	The required 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.

C.2.1.2 Definition

The TOSCA Credential type is defined as follows:

```
tosca.datatypes.Credential:
  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
```

1905 **C.2.1.3 Additional requirements**

- 1906
- TOSCA Orchestrators SHALL interpret and validate the value of the **token** property based upon the value of the **token_type** property.
- 1907

1908 **C.2.1.4 Notes**

- 1909
- Specific token types and encoding them using network protocols are not defined or covered in this specification.
- 1910
- The use of transparent user names (IDs) or passwords are not considered best practice.
- 1911

1912 **C.2.1.5 Examples**

1913 ***C.2.1.5.1 Provide a simple user name and password without a protocol or standardized token format***

```
<some_tosca_entity>:
  properties:
    my_credential:
      type: Credential
      properties:
        user: myusername
        token: mypassword
```

1914 ***C.2.1.5.2 HTTP Basic access authentication credential***

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

```

1915 **C.2.1.5.3 X-Auth-Token credential**

```

<some_tosca_entity>:
  properties:
    my_credential:
      type: Credential
      properties:
        protocol: xauth
        token_type: X-Auth-Token
        # token encoded in Base64
        token: 604bbe45ac7143a79e14f3158df67091

```

1916 **C.2.1.5.4 OAuth bearer token credential**

```

<some_tosca_entity>:
  properties:
    my_credential:
      type: Credential
      properties:
        protocol: oauth2
        token_type: bearer
        # token encoded in Base64
        token: 8ao9nE2DEjr1zCsicWMpBC

```

1917 **C.2.2 tosca.datatypes.network.NetworkInfo**

1918 The Network type is a complex TOSCA data type used to describe logical network information.

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

1919 **C.2.2.1 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.

1920 C.2.2.2 Definition

1921 The TOSCA NetworkInfo data type is defined as follows:

```
tosca.datatypes.network.NetworkInfo:
  properties:
    network_name:
      type: string
    network_id:
      type: string
    addresses:
      type: list
    entry_schema:
      type: string
```

1922 C.2.2.3 Examples

1923 Example usage of the NetworkInfo data type:

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

1924 C.2.2.4 Additional Requirements

- 1925 • It is expected that TOSCA orchestrators MUST be able to map the **network_name** from the TOSCA model
- 1926 to underlying network model of the provider.
- 1927 • The properties (or attributes) of NetworkInfo may or may not be required depending on usage context.

1928 C.2.3 tosca.datatypes.network.PortInfo

1929 The PortInfo type is a complex TOSCA data type used to describe network port information.

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

1930 C.2.3.1 Properties

Name	Type	Constraints	Description
port_name	string	None	The logical network port name.
port_id	string	None	The unique ID for the network port generated by the network provider.
network_id	string	None	The unique ID for the network.
mac_address	string	None	The unique media access control address (MAC address) assigned to the port.

Name	Type	Constraints	Description
addresses	string []	None	The list of IP address(es) assigned to the port.

1931 **C.2.3.2 Definition**

1932 The TOSCA Port type is defined as follows:

```
tosca.datatypes.network.PortInfo:
  properties:
    port_name:
      type: string
    port_id:
      type: string
    network_id:
      type: string
    mac_address:
      type: string
    addresses:
      type: list
    entry_schema:
      type: string
```

1933 **C.2.3.3 Examples**

1934 Example usage of the PortInfo data type:

```
ethernet_port:
  port_name: port1
  port_id: 2c0c7a37-691a-23a6-7709-2d10ad041467
  network_id: 3e54214f-5c09-1bc9-9999-44100326da1b
  mac_address: f1:18:3b:41:92:1e
  addresses: [ 172.24.9.102 ]
```

1935 **C.2.3.4 Additional Requirements**

- 1936 • It is expected that TOSCA orchestrators MUST be able to map the **port_name** from the TOSCA model to
- 1937 underlying network model of the provider.
- 1938 • The properties (or attributes) of PortInfo may or may not be required depending on usage context.

1939 **C.2.4 tosca.datatypes.network.PortDef**

1940 The PortDef type is a TOSCA data Type used to define a network port.

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

1941 **C.2.4.1 Definition**

1942 The TOSCA PortDef type is defined as follows:

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

1943 **C.2.4.2 Examples**

1944 Example use of a PortDef property type:

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

1945 **C.2.5 tosca.datatypes.network.PortSpec**

1946 The PortSpec type is a complex TOSCA data Type used when describing port specifications for a network
1947 connection.

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

1948 **C.2.5.1 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.

1949 **C.2.5.2 Definition**

1950 The TOSCA PortSpec type is defined as follows:

```
tosca.datatypes.network.PortSpec:
```

```

properties:
  protocol:
    type: string
    required: true
    default: tcp
    constraints:
      - valid_values: [ udp, tcp, igmp ]
  target:
    type: integer
    entry_schema:
      type: PortDef
  target_range:
    type: range
    constraints:
      - in_range: [ 1, 65535 ]
  source:
    type: integer
    entry_schema:
      type: PortDef
  source_range:
    type: range
    constraints:
      - in_range: [ 1, 65535 ]

```

1951 C.2.5.3 Additional requirements

- 1952
- A valid PortSpec must have at least one of the following properties: **target**, **target_range**, **source** or **source_range**.
- 1953

1954 C.2.5.4 Examples

1955 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 ]

```

1956 C.3 Capabilities Types

1957 C.3.1 `tosca.capabilities.Root`

1958 This is the default (root) TOSCA Capability Type definition that all other TOSCA Capability Types derive from.

1959 C.3.1.1 Definition

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

1960 C.3.2 `tosca.capabilities.Node`

1961 The Node capability indicates the base capabilities of a TOSCA Node Type.

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

1962 C.3.2.1 Definition

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

1963 C.3.3 `tosca.capabilities.Container`

1964 The Container capability, when included on a Node Type or Template definition, indicates that the node can act
1965 as a container for (or a host for) one or more other declared Node Types.

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

1966 C.3.3.1 Properties

Name	Required	Type	Constraints	Description
num_cpus	no	integer	greater_or_equal: 1	Number of (actual or virtual) CPUs associated with the Compute node.
cpu_frequency	no	scalar- unit.frequency	greater_or_equal: 0.1 GHz	Specifies the operating frequency of CPU's core. This property expresses the expected frequency of one (1) CPU as provided by the property " num_cpus ".
disk_size	no	scalar- unit.size	greater_or_equal: 0 MB	Size of the local disk available to applications running on the Compute node (default unit is MB).
mem_size	no	scalar- unit.size	greater_or_equal: 0 MB	Size of memory available to applications running on the Compute node (default unit is MB).

1967

1968

C.3.3.2 Definition

```

tosca.capabilities.Container:
  derived_from: tosca.capabilities.Root
  properties:
    num_cpus:
      type: integer
      required: false
      constraints:
        - greater_or_equal: 1
    cpu_frequency:
      type: scalar-unit.frequency
      required: false
      constraints:
        - greater_or_equal: 0.1 GHz
    disk_size:
      type: scalar-unit.size
      required: false
      constraints:
        - greater_or_equal: 0 MB
    mem_size:
      type: scalar-unit.size
      required: false
      constraints:
        - greater_or_equal: 0 MB

```

1969

C.3.4 [tosca.capabilities.Endpoint](#)

1970

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

1971

1972

1973

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

1974

C.3.4.1 Properties

Name	Required	Type	Constraints	Description
protocol	yes	string	default: tcp	The name of the protocol (i.e., the protocol prefix) that the endpoint accepts (any OSI Layer 4-7 protocols) Examples: http, https, ftp, tcp, udp, etc.

Name	Required	Type	Constraints	Description
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)

1975 C.3.4.2 Attributes

Name	Required	Type	Constraints	Description
ip_address	yes	string	None	Note: This is the IP address as propagated up by the associated node's host (Compute) container.

1976 C.3.4.3 Definition

```

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

```

```

    type: string
    required: false
  network_name:
    type: string
    required: false
    default: PRIVATE
  initiator:
    type: string
    default: source
    constraints:
      - valid_values: [ source, target, peer ]
  ports:
    type: map
    required: false
    constraints:
      - min_length: 1
    entry_schema:
      type: PortSpec
  attributes:
    ip_address:
      type: string

```

1977 **C.3.4.4 Additional requirements**

- 1978
 - Although both the port and ports properties are not required, one of port or ports must be provided in a
- 1979 valid [Endpoint](#).

1980 **C.3.5 tosca.capabilities.Endpoint.Public**

1981 This capability represents a public endpoint which is accessible to the general internet (and its public IP address

1982 ranges).

1983 This public endpoint capability also can be used to create a floating (IP) address that the underlying network

1984 assigns from a pool allocated from the application’s underlying public network. This floating address is managed

1985 by the underlying network such that can be routed an application’s private address and remains reliable to

1986 internet clients.

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

1987 **C.3.5.1 Definition**

```

tosca.capabilities.Endpoint.Public:
  derived_from: tosca.capabilities.Endpoint

```

```

properties:
  # Change the default network_name to use the first public network found
  network_name: PUBLIC
  floating:
    description: >
      indicates that the public address should be allocated from a pool of
      floating IPs that are associated with the network.
    type: boolean
    default: false
    status: experimental
  dns_name:
    description: The optional name to register with DNS
    type: string
    required: false
    status: experimental

```

1988 **C.3.5.2 Additional requirements**

- 1989
- 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.
- 1990
- If a **dns_name** is set, TOSCA Orchestrators SHALL attempt to register the name in the (local) DNS registry for the Cloud provider.
- 1991
- 1992
- 1993

1994 **C.3.6 tosca.capabilities.Endpoint.Admin**

1995 This is the default TOSCA type that should be used or extended to define a specialized administrator endpoint capability.

1996

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

1997 **C.3.6.1 Properties**

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

1998 **C.3.6.2 Definition**

```

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

```

1999 **C.3.6.3 Additional requirements**

- 2000 • TOSCA Orchestrator implementations of Endpoint.Admin (and connections to it) **SHALL** assure that
- 2001 network-level security is enforced if possible.

2002 **C.3.7 tosca.capabilities.Endpoint.Database**

2003 This is the default TOSCA type that should be used or extended to define a specialized database endpoint
2004 capability.

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

2005 **C.3.7.1 Properties**

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

2006 **C.3.7.2 Definition**

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

2007 **C.3.8 tosca.capabilities.Attachment**

2008 This is the default TOSCA type that should be used or extended to define an attachment capability of a (logical)
2009 infrastructure device node (e.g., [BlockStorage](#) node).

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

2010 **C.3.8.1 Properties**

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

2011 **C.3.8.2 Definition**

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

2012 **C.3.9 tosca.capabilities.OperatingSystem**

2013 This is the default TOSCA type that should be used to express an Operating System capability for a node.

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

2014

C.3.9.1 Properties

Name	Required	Type	Constraints	Description
architecture	no	string	None	The Operating System (OS) architecture. Examples of valid values include: x86_32, x86_64, etc.
type	no	string	None	The Operating System (OS) type. Examples of valid values include: linux, aix, mac, windows, etc.
distribution	no	string	None	The Operating System (OS) distribution. Examples of valid values for an “type” of “Linux” would include: debian, fedora, rhel and ubuntu.
version	no	version	None	The Operating System version.

2015

C.3.9.2 Definition

```
tosca.capabilities.OperatingSystem:
  derived_from: tosca.capabilities.Root
  properties:
    architecture:
      type: string
      required: false
    type:
      type: string
      required: false
    distribution:
      type: string
      required: false
    version:
      type: version
      required: false
```

2016

C.3.9.3 Additional Requirements

2017
2018
2019
2020

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

2021 **C.3.9.4 Notes**

- 2022
 - None

2023 **C.3.10 tosca.capabilities.Scalable**

2024 This is the default TOSCA type that should be used to express a scalability capability for a node.

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

2025 **C.3.10.1 Properties**

Name	Required	Type	Constraints	Description
min_instances	yes	integer	default: 1	This property is used to indicate the minimum number of instances that should be created for the associated TOSCA Node Template by a TOSCA orchestrator.
max_instances	yes	integer	default: 1	This property is used to indicate the maximum number of instances that should be created for the associated TOSCA Node Template by a TOSCA orchestrator.
default_instances	no	integer	N/A	An optional property that indicates the requested default number of instances that should be the starting number of instances a TOSCA orchestrator should attempt to allocate. Note: The value for this property MUST be in the range between the values set for 'min_instances' and 'max_instances' properties.

2026 **C.3.10.2 Definition**

```

tosca.capabilities.Scalable:
  derived_from: tosca.capabilities.Root
  properties:
    min_instances:
      type: integer
      default: 1
    max_instances:
      type: integer
      default: 1
    default_instances:
      type: integer

```

2027 **C.3.10.3 Notes**

- 2028
 - The actual number of instances for a node may be governed by a separate scaling policy which
- 2029 conceptually would be associated to either a scaling-capable node or a group of nodes in which it is
- 2030 defined to be a part of. This is a planned future feature of the TOSCA Simple Profile and not currently
- 2031 described.

2032 **C.3.11 tosca.capabilities.network.Bindable**

2033 A node type that includes the Bindable capability indicates that it can be bound to a logical network association
2034 via a network port.

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

2035 **C.3.11.1 Properties**

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

2036 **C.3.11.2 Definition**

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

2037 **C.4 Requirement Types**

2038 There are no normative Requirement Types currently defined in this working draft. Typically, Requirements are
2039 described against a known Capability Type

2040 **C.5 Relationship Types**

2041 **C.5.1 tosca.relationships.Root**

2042 This is the default (root) TOSCA Relationship Type definition that all other TOSCA Relationship Types derive
2043 from.

2044 **C.5.1.1 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.

2045

2046 **C.5.1.2 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

```

2047 **C.5.2 tosca.relationships.DependsOn**

2048 This type represents a general dependency relationship between two nodes.

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

2049 **C.5.2.1 Definition**

```

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

```

2050 **C.5.3 tosca.relationships.HostedOn**

2051 This type represents a hosting relationship between two nodes.

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

2052 **C.5.3.1 Definition**

```

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

```

2053 **C.5.4 tosca.relationships.ConnectsTo**

2054 This type represents a network connection relationship between two nodes.

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

2055 **C.5.4.1 Definition**

```
tosca.relationships.ConnectsTo:
  derived_from: tosca.relationships.Root
  valid_target_types: [ tosca.capabilities.Endpoint ]
  properties:
    credential:
      type: tosca.datatypes.Credential
      required: false
```

2056 **C.5.4.2 Properties**

Name	Required	Type	Constraints	Description
credential	no	Credential	None	The security credential to use to present to the target endpoint to for either authentication or authorization purposes.

2057 **C.5.5 tosca.relationships.AttachesTo**

2058 This type represents an attachment relationship between two nodes. For example, an AttachesTo relationship
 2059 type would be used for attaching a storage node to a Compute node.

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

2060 **C.5.5.1 Properties**

Name	Required	Type	Constraints	Description
location	yes	string	min_length: 1	The relative location (e.g., path on the file system), which provides the root location to address an attached node. e.g., a mount point / path such as '/usr/data' Note: The user must provide it and it cannot be "root".
device	no	string	None	The logical device name which for the attached device (which is represented by the target node in the model). e.g., '/dev/hda1'

2061 **C.5.5.2 Attributes**

Name	Required	Type	Constraints	Description
device	no	string	None	The logical name of the device as exposed to the instance. Note: A runtime property that gets set when the model gets instantiated by the orchestrator.

2062 **C.5.5.3 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

```

2063 **C.5.6 tosca.relationships.RoutesTo**

2064 This type represents an intentional network routing between two Endpoints in different networks.

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

2065 **C.5.6.1 Properties**

Name	Required	Type	Constraints	Description
TBD				

2066 **C.5.6.2 Attributes**

Name	Required	Type	Constraints	Description
TBD				

2067 **C.5.6.3 Definition**

```

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

```

2068 C.6 Interface Types

2069 Interfaces are reusable entities that define a set of operations that that can be included as part of a Node type or
2070 Relationship Type definition. Each named operations may have code or scripts associated with them that
2071 orchestrators can execute for when transitioning an application to a given state.

2072 C.6.1 Additional Requirements

- 2073 • Designers of Node or Relationship types are not required to actually provide/associate code or scripts
2074 with every operation for a given interface it supports. In these cases, orchestrators SHALL consider that
2075 a “No Operation” or “no-op”.
- 2076 • The default behavior when providing scripts for an operation in a sub-type (sub-class) or a template of
2077 an existing type which already has a script provided for that operation SHALL be override. Meaning that
2078 the subclasses’ script is used in place of the parent type’s script.

2079 C.6.2 Best Practices

- 2080 • When TOSCA Orchestrators substitute an implementation for an abstract node in a deployed service
2081 template it SHOULD be able to present a confirmation to the submitter to confirm the implementation
2082 chosen would be acceptable.

2083 C.6.3 `tosca.interfaces.node.lifecycle.Standard`

2084 This lifecycle interface defines the essential, normative operations that TOSCA nodes may support.

Shorthand Name	Standard
Type Qualified Name	tosca: Standard
Type URI	tosca.interfaces.node.lifecycle.Standard

2085 C.6.3.1 Definition

```
tosca.interfaces.node.lifecycle.Standard:  
  create:  
    description: Standard lifecycle create operation.  
  configure:  
    description: Standard lifecycle configure operation.  
  start:  
    description: Standard lifecycle start operation.  
  stop:  
    description: Standard lifecycle stop operation.  
  delete:  
    description: Standard lifecycle delete operation.
```

2086 C.6.3.2 Create operation

2087 The create operation is generally used to create the resource or service the node represents in the topology.
2088 TOSCA orchestrators expect node templates to provide either a deployment artifact or an implementation artifact
2089 of a defined artifact type that it is able to process. This specification defines normative deployment and

2090 implementation artifact types all TOSCA Orchestrators are expected to be able to process to support application
2091 portability.

2092

2093 C.6.3.3 Using TOSCA artifacts with interface operations

2094 # Bash script whose location is described in the TOSCA CSAR file

2095 # 1) LOCAL SCOPE: symbolic artifact name <or> 2) GLOBAL: actual file name from CSAR

2096 #

2097 C.6.3.4 TOSCA Orchestrator processing of Deployment artifacts

2098 TOSCA Orchestrators, when encountering a deployment artifact on the create operation; will automatically
2099 attempt to deploy the artifact based upon its artifact type. This means that no implementation artifacts (e.g.,
2100 scripts) are needed on the create operation to provide commands that deploy or install the software.

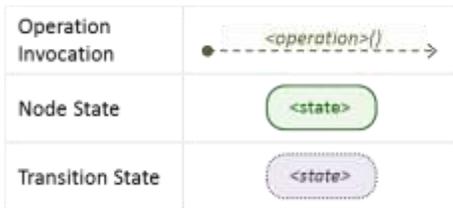
2101

2102 For example, if a TOSCA Orchestrator is processing an application with a node of type SoftwareComponent and
2103 finds that the node's template has a create operation that provides a filename (or references to an artifact which
2104 describes a file) of a known TOSCA deployment artifact type such as an Open Virtualization Format (OVF) image
2105 it will automatically deploy that image into the SoftwareComponent's host Compute node.

2106 C.6.3.5 Operation sequencing and node state

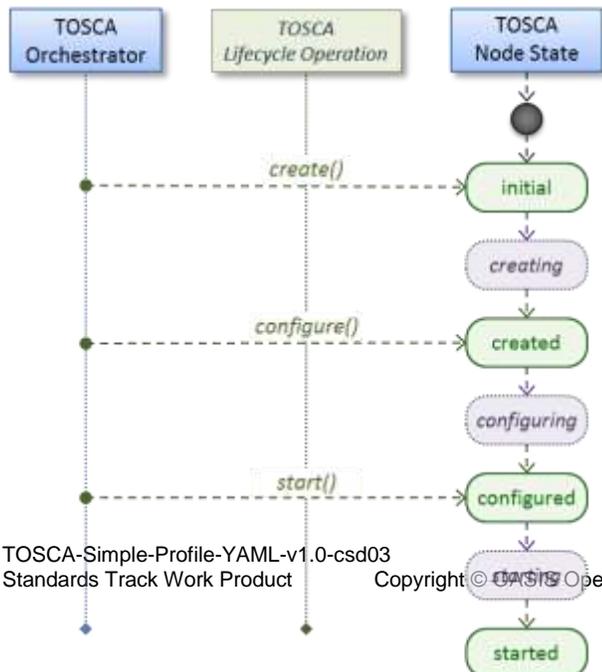
2107 The following diagrams show how TOSCA orchestrators sequence the operations of the Standard lifecycle in
2108 normal node startup and shutdown procedures.

2109 The following key should be used to interpret the diagrams:



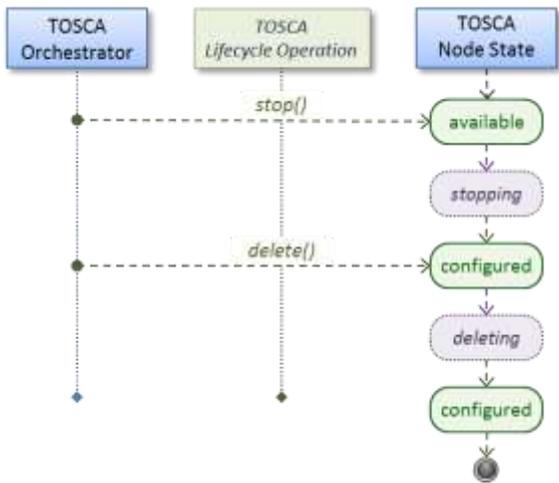
2110 C.6.3.5.1 Normal node startup sequence diagram

2111 The following diagram shows how the TOSCA orchestrator would invoke operations on the Standard lifecycle to
2112 shut down a node.



2113 **C.6.3.5.2 Normal node shutdown sequence diagram**

2114 The following diagram shows how the TOSCA orchestrator would invoke operations on the Standard lifecycle to
 2115 shut down a node.



2116 **C.6.4 tosca.interfaces.relationship.Configure**

2117 The lifecycle interfaces define the essential, normative operations that each TOSCA Relationship Types may
 2118 support.

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

2119 **C.6.4.1 Definition**

```

tosca.interfaces.relationship.Configure:
  pre_configure_source:
    description: Operation to pre-configure the source endpoint.
  pre_configure_target:
    description: Operation to pre-configure the target endpoint.
  post_configure_source:
    description: Operation to post-configure the source endpoint.
  post_configure_target:
    description: Operation to post-configure the target endpoint.
  add_target:
    description: Operation to notify the source node of a target node being added
    via a relationship.
  add_source:
    description: Operation to notify the target node of a source node which is now
    available via a relationship.
  description:
  target_changed:
    description: Operation to notify source some property or attribute of the target
  
```

changed

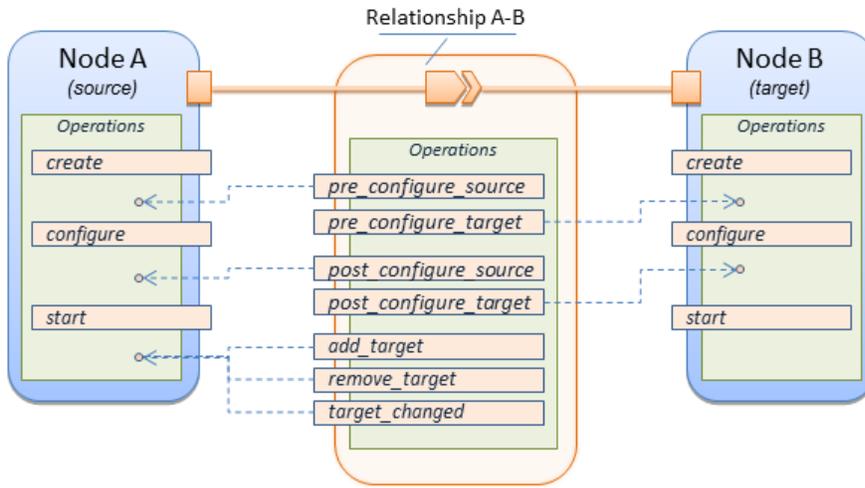
remove_target:

description: Operation to remove a target node.

2120

2121

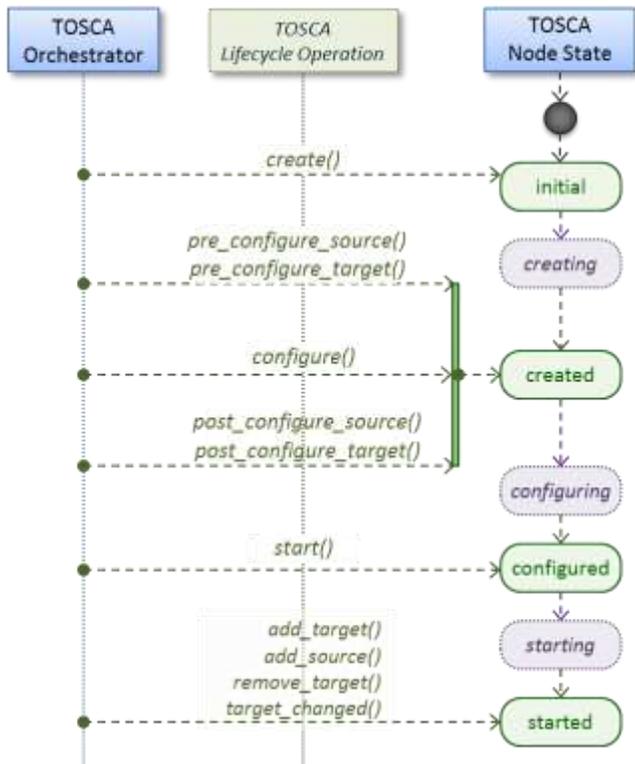
2122 C.6.4.2 Invocation Conventions



2123 TOSCA relationships are directional connecting a source node to a target node. When TOSCA Orchestrator
 2124 connects a source and target node together using a relationship that supports the Configure interface it will
 2125 “interleave” the operations invocations of the Configure interface with those of the node’s own Standard lifecycle
 2126 interface. This concept is illustrated below:

2127 C.6.4.3 Normal node start sequence with Configure relationship operations

2128 The following diagram shows how the TOSCA orchestrator would invoke Configure lifecycle operations in
 2129 conjunction with Standard lifecycle operations during a typical startup sequence on a node.



2130 **C.6.4.3.1 Node-Relationship configuration sequence**

2131 Depending on which side (i.e., source or target) of a relationship a node is on, the orchestrator will:

2132 Invoke either the **pre_configure_source** or **pre_configure_target** operation as supplied by the
2133 relationship on the node.

2134 Invoke the node's **configure** operation.

2135 Invoke either the **post_configure_source** or **post_configure_target** as supplied by the relationship
2136 on the node.

2137 Note that the **pre_configure_xxx** and **post_configure_xxx** are invoked only once per node instance.

2138 **C.6.4.3.2 Node-Relationship add, remove and changed sequence**

2139 Since a topology template contains nodes that can dynamically be added (and scaled), removed or changed as
2140 part of an application instance, the Configure lifecycle includes operations that are invoked on node instances that
2141 to notify and address these dynamic changes.

2142

2143 For example, a source node, of a relationship that uses the Configure lifecycle, will have the relationship
2144 operations **add_target**, or **remove_target** invoked on it whenever a target node instance is added or removed
2145 to the running application instance. In addition, whenever the node state of its target node changes, the
2146 **target_changed** operation is invoked on it to address this change. Conversely, the **add_source** and
2147 **remove_source** operations are invoked on the source node of the relationship.

2148 **C.6.4.4 Notes**

- 2149 • The target (provider) MUST be active and running (i.e., all its dependency stack MUST be fulfilled) prior
2150 to invoking **add_target**
- 2151 • In other words, all Requirements MUST be satisfied before it advertises its capabilities (i.e., the
2152 attributes of the matched Capabilities are available).
- 2153 • In other words, it cannot be “consumed” by any dependent node.
- 2154 • Conversely, since the source (consumer) needs information (attributes) about any targets (and their
2155 attributes) being removed before it actually goes away.
- 2156 • The **remove_target** operation should only be executed if the target has had **add_target** executed. BUT
2157 in truth we're first informed about a target in **pre_configure_source**, so if we execute that the source
2158 node should see **remove_target** called to cleanup.
- 2159 • **Error handling:** If any node operation of the topology fails processing should stop on that node template
2160 and the failing operation (script) should return an error (failure) code when possible.

2161 **C.7 Node Types**

2162 **C.7.1 toscanodes.Root**

2163 The TOSCA **Root** Node Type is the default type that all other TOSCA base Node Types derive from. This allows
2164 for all TOSCA nodes to have a consistent set of features for modeling and management (e.g., consistent
2165 definitions for requirements, capabilities and lifecycle interfaces).

2166

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

2167

C.7.1.1 Properties

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

2168

C.7.1.2 Attributes

Name	Required	Type	Constraints	Description
tosca_id	yes	string	None	A unique identifier of the realized instance of a Node Template that derives from any TOSCA normative type.
tosca_name	yes	string	None	This attribute reflects the name of the Node Template as defined in the TOSCA service template. This name is not unique to the realized instance model of corresponding deployed application as each template in the model can result in one or more instances (e.g., scaled) when orchestrated to a provider environment.
state	yes	string	default: initial	The state of the node instance. See section “Node States” for allowed values.

2169

C.7.1.3 Definition

```

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

```

type: [tosca.interfaces.node.lifecycle.Standard](#)

2170 C.7.1.4 Additional Requirements

- 2171 • All Node Type definitions that wish to adhere to the TOSCA Simple Profile **SHOULD** extend from the TOSCA
2172 Root Node Type to be assured of compatibility and portability across implementations.

2173 C.7.2 [tosca.nodes.Compute](#)

2174 The TOSCA **Compute** node represents one or more real or virtual processors of software applications or services
2175 along with other essential local resources. Collectively, the resources the compute node represents can logically
2176 be viewed as a (real or virtual) “server”.

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

2177 C.7.2.1 Properties

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

2178 C.7.2.2 Attributes

Name	Required	Type	Constraints	Description
private_address	no	string	None	The primary private IP address assigned by the cloud provider that applications may use to access the Compute node.
public_address	no	string	None	The primary public IP address assigned by the cloud provider that applications may use to access the Compute node.
networks	no	map of NetworkInfo	None	The list of logical networks assigned to the compute host instance and information about them.
ports	no	map of PortInfo	None	The list of logical ports assigned to the compute host instance and information about them.

2179 C.7.2.3 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

```

2180 C.7.3 [tosca.nodes.SoftwareComponent](#)

2181 The TOSCA **SoftwareComponent** node represents a generic software component that can be managed and run
 2182 by a TOSCA **Compute** Node Type.

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

2183 C.7.3.1 Properties

Name	Required	Type	Constraints	Description
component_version	no	version	None	The software component's version.

2184 C.7.3.2 Attributes

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

C.7.3.3 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

```

C.7.3.4 Additional Requirements

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

C.7.4 [tosca.nodes.WebServer](#)

This TOSA **WebServer** Node Type represents an abstract software component or service that is capable of hosting and providing management operations for one or more **WebApplication** nodes.

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

C.7.4.1 Properties

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

2193 **C.7.4.2 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 ]

```

2194 **C.7.4.3 Notes and Additional Requirements**

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

2197 **C.7.5 tosca.nodes.WebApplication**

2198 The TOSCA **WebApplication** node represents a software application that can be managed and run by a TOSCA
 2199 **WebServer** node. Specific types of web applications such as Java, etc. could be derived from this type.

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

2200 **C.7.5.1 Properties**

Name	Required	Type	Constraints	Description
context_root	no	string	None	The web application's context root which designates the application's URL path within the web server it is hosted on.

2201 **C.7.5.2 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
```

2202 C.7.5.3 Additional Requirements

- 2203 • None

2204 C.7.6 [tosca.nodes.DBMS](#)

2205 The TOSCA **DBMS** node represents a typical relational, SQL Database Management System software component
2206 or service.

2207 C.7.6.1 Properties

Name	Required	Type	Constraints	Description
root_password	no	string	None	The optional root password for the DBMS server.
port	no	integer	None	The DBMS server's port.

2208 C.7.6.2 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 ]
```

2209 C.7.6.3 Additional Requirements

- 2210 • None

2211 C.7.7 [tosca.nodes.Database](#)

2212 The TOSCA **Database** node represents a logical database that can be managed and hosted by a TOSCA **DBMS**
2213 node.

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

2214

C.7.7.1 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.

2215

C.7.7.2 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

```

2216 **C.7.8 tosca.nodes.ObjectStorage**

2217 The TOSCA **ObjectStorage** node represents storage that provides the ability to store data as objects (or BLOBs
2218 of data) without consideration for the underlying filesystem or devices.

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

2219 **C.7.8.1 Properties**

Name	Required	Type	Constraints	Description
name	yes	string	None	The logical name of the object store (or container).
size	no	scalar-unit.size	greater_or_equal: 0 GB	The requested initial storage size (default unit is in Gigabytes).
maxsize	no	scalar-unit.size	greater_or_equal: 0 GB	The requested maximum storage size (default unit is in Gigabytes).

2220 **C.7.8.2 Definition**

```

tosca.nodes.ObjectStorage:
  derived_from: tosca.nodes.Root
  properties:
    name:
      type: string
    size:
      type: scalar-unit.size
      constraints:
        - greater_or_equal: 0 GB
    maxsize:
      type: scalar-unit.size
      constraints:
        - greater_or_equal: 0 GB
  capabilities:
    storage_endpoint:
      type: tosca.capabilities.Endpoint

```

2221 **C.7.8.3 Notes:**

- 2222 • Subclasses of the ObjectStorage node may impose further constraints on properties. For example, a
2223 subclass may constrain the (minimum or maximum) length of the 'name' property or include a regular
2224 expression to constrain allowed characters used in the 'name' property.

2225 **C.7.9 tosca.nodes.BlockStorage**

2226 The TOSCA **BlockStorage** node currently represents a server-local block storage device (i.e., not shared)
2227 offering evenly sized blocks of data from which raw storage volumes can be created.

2228 **Note:** In this draft of the TOSCA Simple Profile, distributed or Network Attached Storage (NAS) are not yet
 2229 considered (nor are clustered file systems), but the TC plans to do so in future drafts.

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

2230 **C.7.9.1 Properties**

Name	Required	Type	Constraints	Description
size	yes *	scalar-unit.size	greater_or_equal: 1 MB	The requested storage size (default unit is MB). * Note: <ul style="list-style-type: none"> • Required when an existing volume (i.e., volume_id) is not available. • If volume_id is provided, size is ignored. Resize of existing volumes is not considered at this time.
volume_id	no	string	None	ID of an existing volume (that is in the accessible scope of the requesting application).
snapshot_id	no	string	None	Some identifier that represents an existing snapshot that should be used when creating the block storage (volume).

2231 **C.7.9.2 Attributes**

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

2232 **C.7.9.3 Definition**

```
tosca.nodes.BlockStorage:
  derived_from: tosca.nodes.Root
  properties:
    size:
      type: scalar-unit.size
      constraints:
        - greater_or_equal: 1 MB
    volume_id:
      type: string
      required: false
    snapshot_id:
      type: string
      required: false
  capabilities:
    attachment:
      type: tosca.capabilities.Attachment
```

2233 C.7.9.4 Additional Requirements

- 2234 • The **size** property is required when an existing volume (i.e., **volume_id**) is not available. However, if the
2235 property **volume_id** is provided, the **size** property is ignored.

2236 C.7.9.5 Notes

- 2237 • Resize is of existing volumes is not considered at this time.
- 2238 • It is assumed that the volume contains a single filesystem that the operating system (that is hosting an
2239 associate application) can recognize and mount without additional information (i.e., it is operating
2240 system independent).
- 2241 • Currently, this version of the Simple Profile does not consider regions (or availability zones) when
2242 modeling storage.

2243 C.7.10 **tosca.nodes.Container.Runtime**

2244 The TOSCA **Container** Runtime node represents operating system-level virtualization technology used to run
2245 multiple application services on a single Compute host.

Shorthand Name	Container.Runtime
Type Qualified Name	tosca:Container.Runtime
Type URI	tosca.nodes.Container.Runtime

2246 C.7.10.1 Definition

```
tosca.nodes.Container.Runtime:  
  derived_from: tosca.nodes.SoftwareComponent  
  capabilities:  
    host:  
      type: tosca.capabilities.Container  
    scalable:  
      type: tosca.capabilities.Scalable
```

2247 C.7.11 **tosca.nodes.Container.Application**

2248 The TOSCA **Container** Application node represents an application that requires **Container**-level virtualization
2249 technology.

Shorthand Name	Container.Application
Type Qualified Name	tosca:Container.Application
Type URI	tosca.nodes.Container.Application

2250 C.7.11.1 Definition

```
tosca.nodes.Container.Application:  
  derived_from: tosca.nodes.Root  
  requirements:
```

```

- host:
  capability: tosca.capabilities.Container
  node: tosca.nodes.Container
  relationship: tosca.relationships.HostedOn

```

2251 C.7.12 [tosca.nodes.LoadBalancer](#)

2252 The TOSCA **Load Balancer** node represents logical function that be used in conjunction with a Floating Address
 2253 to distribute an application’s traffic (load) across a number of instances of the application (e.g., for a clustered or
 2254 scaled application).

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

2255 C.7.12.1 Definition

```

tosca.nodes.LoadBalancer:
  derived_from: tosca.nodes.Root
  properties:
    # TBD
  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

```

2256 C.7.13 Notes:

- 2257 • A **LoadBalancer** node can still be instantiated and managed independently of any applications it would
 2258 serve; therefore, the load balancer’s **application** requirement allows for zero occurrences.

2259 C.8 Artifact Types

2260 TOSCA Artifacts represent the packages and imperative used by the orchestrator when invoking TOSCA
 2261 Interfaces on Node or Relationship Types. Currently, artifacts are logically divided into three categories:

2262
2263
2264
2265
2266
2267
2268
2269
2270
2271
2272

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

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

C.8.1 `tosca.artifacts.Root`

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

C.8.1.1 Definition

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

2276
2277
2278
2279

C.8.2 `tosca.artifacts.File`

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

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

C.8.2.1 Definition

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

2281
2282
2283
2284
2285
2286

C.8.3 Deployment Types

C.8.3.1 `tosca.artifacts.Deployment`

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

C.8.3.1.1 Definition

```
tosca.artifacts.Deployment:
```

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

2287 **C.8.3.2 [tosca.artifacts.Deployment.Image](#)**

2288 This artifact type represents a parent type for any “image” which is an opaque packaging of a TOSCA Node’s
2289 deployment (whether real or virtual) whose contents are typically already installed and pre-configured (i.e.,
2290 “stateful”) and prepared to be run on a known target container.

Shorthand Name	Deployment.Image
Type Qualified Name	tosca:Deployment.Image
Type URI	tosca.artifacts.Deployment.Image

2291 **C.8.3.2.1 Definition**

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

2292 **C.8.4 [tosca.artifacts.Deployment.Image.VM](#)**

2293 This artifact represents the parent type for all Virtual Machine (VM) image and container formatted deployment
2294 artifacts. These images contain a stateful capture of a machine (e.g., server) including operating system and
2295 installed software along with any configurations and can be run on another machine using a hypervisor which
2296 virtualizes typical server (i.e., hardware) resources.

2297 **C.8.4.1 Definition**

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

2298 **C.8.4.2 Notes:**

- 2299 • Future drafts of this specification may include popular standard VM disk image (e.g., ISO, VMI, VMDX,
2300 QCOW2, etc.) and container (e.g., OVF, bare, etc.) formats. These would include consideration of disk
2301 formats such as:

2302 **C.8.5 Implementation Types**

2303 **C.8.5.1 [tosca.artifacts.Implementation](#)**

2304 This artifact type represents the parent type for all implementation artifacts in TOSCA. These artifacts are used
2305 to implement operations of TOSCA interfaces either directly (e.g., scripts) or indirectly (e.g., config. files).

2306 **C.8.5.1.1 Definition**

```
tosca.artifacts.Implementation:
  derived_from: tosca.artifacts.Root
```

description: TOSCA base type for implementation artifacts

2307 **C.8.5.2 `tosca.artifacts.Implementation.Bash`**

2308 This artifact type represents a Bash script type that contains Bash commands that can be executed on the Unix
2309 Bash shell.

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

2310 **C.8.5.2.1 Definition**

```
tosca.artifacts.Implementation.Bash:  
  derived_from: toasca.artifacts.Implementation  
  description: Script artifact for the Unix Bash shell  
  mime_type: application/x-sh  
  file_ext: [ sh ]
```

2311 **C.8.5.3 `tosca.artifacts.Implementation.Python`**

2312 This artifact type represents a Python file that contains Python language constructs that can be executed within
2313 a Python interpreter.

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

2314 **C.8.5.3.1 Definition**

```
tosca.artifacts.Implementation.Python:  
  derived_from: toasca.artifacts.Implementation  
  description: Artifact for the interpreted Python language  
  mime_type: application/x-python  
  file_ext: [ py ]
```

2315 Appendix D. Non-normative type definitions

2316 This section defines non-normative types used in examples or use cases within this specification.

2317 D.1 Artifact Types

2318 D.1.1 `tosca.artifacts.Deployment.Image.Container.Docker`

2319 This artifact represents a Docker “image” (a TOSCA deployment artifact type) which is a binary comprised of one
2320 or more (a union of read-only and read-write) layers created from snapshots within the underlying Docker **Union**
2321 **File System**.

2322 D.1.1.1 Definition

```
tosca.artifacts.Deployment.Image.Container.Docker:  
  derived_from: tosca.artifacts.Deployment.Image  
  description: Docker Container Image
```

2323 D.1.2 `tosca.artifacts.Deployment.Image.VM.ISO`

2324 A Virtual Machine (VM) formatted as an ISO standard disk image.

2325 D.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
```

2326 D.2 Capability Types

2327 D.2.1 `tosca.capabilities.Container.Docker`

2328 The type indicates capabilities of a Docker runtime environment (client).

Shorthand Name	Container.Docker
Type Qualified Name	tosca:Container.Docker
Type URI	tosca.capabilities.Container.Docker

2329 D.2.1.1 Properties

Name	Required	Type	Constraints	Description
version	no	version[]	None	The Docker version capability (i.e., the versions supported by the capability).
publish_all	no	boolean	default: false	Indicates that all ports (ranges) listed in the <i>dockerfile</i> using the EXPOSE keyword be published.
publish_ports	no	list of PortSpec	None	List of ports mappings from source (Docker container) to target (host) ports to publish.

Name	Required	Type	Constraints	Description
expose_ports	no	list of PortSpec	None	List of ports mappings from source (Docker container) to expose to other Docker containers (not accessible outside host).
volumes	no	list of string	None	The <i>dockerfile</i> VOLUME command which is used to enable access from the Docker container to a directory on the host machine.
host_id	no	string	None	The optional identifier of an existing host resource that should be used to run this container on.
volume_id	no	string	None	The optional identifier of an existing storage volume (resource) that should be used to create the container's mount point(s) on.

2330

D.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

```

2331

D.2.1.3 Additional requirements

2332

- When the **expose_ports** property is used, only the **source** and **source_range** properties of [PortSpec](#) SHALL be valid for supplying port numbers or ranges, the **target** and **target_range** properties are ignored.

2333

2334 D.3 Node Types

2335 D.3.1 `tosca.nodes.Database.MySQL`

2336 D.3.1.1 Properties

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

2337 D.3.1.2 Definition

```
tosca.nodes.Database.MySQL:  
  derived_from: tosca.nodes.Database  
  properties:  
    root_password:  
      required: true  
  requirements:  
    - host:  
      node: tosca.nodes.DBMS.MySQL
```

2338 D.3.2 `tosca.nodes.DBMS.MySQL`

2339 D.3.2.1 Properties

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

2340 D.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  
      required: true  
  capabilities:  
    # Further constrain the 'host' capability to only allow MySQL databases  
    host:  
      valid_source_types: [ tosca.nodes.Database.MySQL ]
```

2341 **D.3.3 [tosca.nodes.WebServer.Apache](#)**

2342 **D.3.3.1 Properties**

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

2343 **D.3.3.2 Definition**

```
tosca.nodes.WebServer.Apache:  
  derived_from: tosca.nodes.WebServer
```

2344 **D.3.4 [tosca.nodes.WebApplication.WordPress](#)**

2345 **D.3.4.1 Properties**

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

2346 **D.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
```

2347 **D.3.5 [tosca.nodes.WebServer.Nodejs](#)**

2348 **D.3.5.1 Properties**

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

2349 **D.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

```

2350 **D.3.6 [tosca.nodes.Container.Application.Docker](#)**

2351 **D.3.6.1 Properties**

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

2352 **D.3.6.2 Definition**

```

tosca.nodes.Container.Application.Docker:
  derived_from: tosca.nodes.Container.Application
  requirements:
    - host:
        capability: tosca.capabilities.Container.Docker

```

2353

Appendix E. TOSCA Cloud Service Archive (CSAR) Format

2354
2355
2356
2357

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.

2358

E.1 Overall Structure of a CSAR

2359
2360
2361
2362
2363

A CSAR zip file is required to contain a **TOSCA-Metadata** directory, which in turn contains the **TOSCA.meta** metadata file that provides entry information for a TOSCA orchestrator processing the CSAR file. The CSAR file may contain other directories with arbitrary names and contents. Note that in contrast to the TOSCA 1.0 specification, it is not required to put TOSCA definitions files into a special “Definitions” directory, but definitions YAML files can be placed into any directory within the CSAR file.

2364

E.2 TOSCA Meta File

2365
2366
2367
2368
2369
2370
2371
2372

The **TOSCA.meta** file structure follows the exact same syntax as defined in the TOSCA 1.0 specification. However, it is only required to include *block_0* (see section 16.2 in [TOSCA-1.0]) with the **Entry-Definitions** keyword pointing to a valid TOSCA definitions YAML file that a TOSCA orchestrator should use as entry for parsing the contents of the overall CSAR file. Note that it is not required to explicitly list TOSCA definitions files in subsequent blocks of the **TOSCA.meta** file, but any TOSCA definitions files besides the one denoted by the **Entry-Definitions** keyword can be found by a TOSCA orchestrator by processing respective **imports** statements in the entry definitions file (or in recursively imported files).

2373
2374
2375
2376

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.

2377
2378

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**.

2379

E.2.1 Example

2380

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
```

2381
2382
2383
2384

This **TOSCA.meta** file indicates its simplified TOSCA Simple Profile structure by means of the **CSAR-Version** keyword with value **1.1**. The **Entry-Definitions** keyword points to a TOSCA definitions YAML file with the name **tosca_elk.yaml** which is contained in a directory called **definitions** within the root of the CSAR file.

2385

Appendix F. Networking

2386

This describes how to express and control the application centric network semantics available in TOSCA.

2387

F.1 Networking and Service Template Portability

2388

TOSCA Service Templates are application centric in the sense that they focus on describing application components in terms of their requirements and interrelationships. In order to provide cloud portability, it is important that a TOSCA Service Template avoid cloud specific requirements and details. However, at the same time, TOSCA must provide the expressiveness to control the mapping of software component connectivity to the network constructs of the hosting cloud.

2389

2390

2391

2392

2393

TOSCA Networking takes the following approach.

2394

1. The application component connectivity semantics and expressed in terms of Requirements and Capabilities and the relationships between these. Service Template authors are able to express the interconnectivity requirements of their software components in an abstract, declarative, and thus highly portable manner.
2. The information provided in TOSCA is complete enough for a TOSCA implementation to fulfill the application component network requirements declaratively (i.e., it contains information such as communication initiation and layer 4 port specifications) so that the required network semantics can be realized on arbitrary network infrastructures.
3. TOSCA Networking provides full control of the mapping of software component interconnectivity to the networking constructs of the hosting cloud network independently of the Service Template, providing the required separation between application and network semantics to preserve Service Template portability.
4. Service Template authors have the choice of specifying application component networking requirements in the Service Template or completely separating the application component to network mapping into a separate document. This allows application components with explicit network requirements to express them while allowing users to control the complete mapping for all software components which may not have specific requirements. Usage of these two approaches is possible simultaneously and required to avoid having to re-write components network semantics as arbitrary sets of components are assembled into Service Templates.
5. Defining a set of network semantics which are expressive enough to address the most common application connectivity requirements while avoiding dependencies on specific network technologies and constructs. Service Template authors and cloud providers are able to express unique/non-portable semantics by defining their own specialized network Requirements and Capabilities.

2395

2396

2397

2398

2399

2400

2401

2402

2403

2404

2405

2406

2407

2408

2409

2410

2411

2412

2413

2414

2415

2416

2417

F.2 Connectivity Semantics

2418

TOSCA's application centric approach includes the modeling of network connectivity semantics from an application component connectivity perspective. The basic premise is that applications contain components which need to communicate with other components using one or more endpoints over a network stack such as TCP/IP, where connectivity between two components is expressed as a <source component, source address, source port, target component, target address, target port> tuple. Note that source and target components are added to the traditional 4 tuple to provide the application centric information, mapping the network to the source or target component involved in the connectivity.

2419

2420

2421

2422

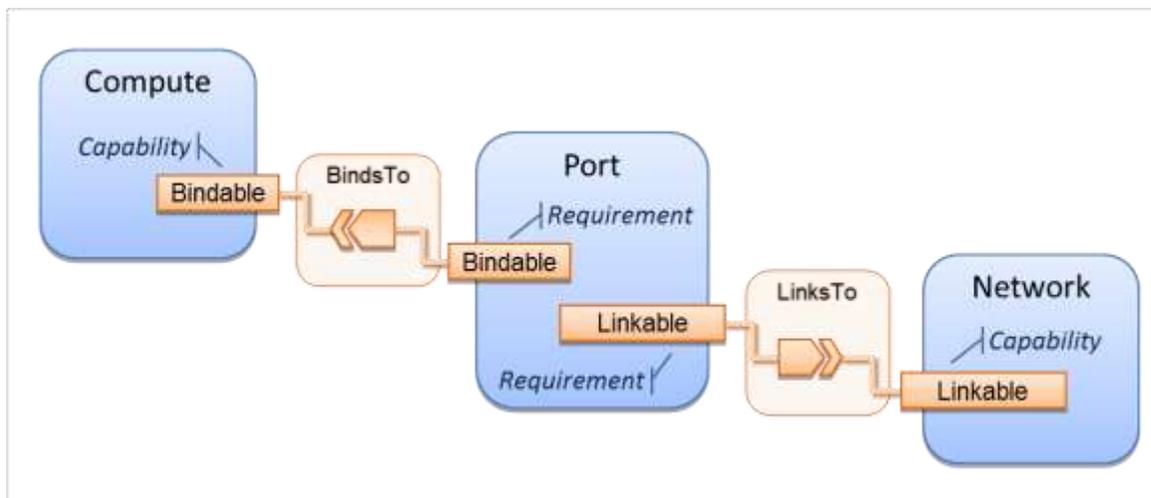
2423

2424

2425

2426 Software components are expressed as Node Types in TOSCA which can express virtually any kind of concept in
2427 a TOSCA model. Node Types offering network based functions can model their connectivity using a special
2428 Endpoint Capability. [tosca.capabilities.Endpoint](#), designed for this purpose. Node Types which require an
2429 Endpoint can specify this as a TOSCA requirement. A special Relationship Type, [tosca.relationships.ConnectsTo](#),
2430 is used to implicitly or explicitly relate the source Node Type's endpoint to the required endpoint in the target node
2431 type. Since [tosca.capabilities.Endpoint](#) and [tosca.relationships.ConnectsTo](#) are TOSCA types, they can be used
2432 in templates and extended by subclassing in the usual ways, thus allowing the expression of additional semantics
2433 as needed.

2434 The following diagram shows how the TOSCA node, capability and relationship types enable modeling the
2435 application layer decoupled from the network model intersecting at the Compute node using the [Bindable](#)
2436 capability type.



2437 As you can see, the Port node type effectively acts a broker node between the Network node description and a
2438 host Compute node of an application.

2439 **F.3 Expressing connectivity semantics**

2440 This section describes how TOSCA supports the typical client/server and group communication semantics found
2441 in application architectures.

2442 **F.3.1 Connection initiation semantics**

2443 The [tosca.relationships.ConnectsTo](#) expresses that requirement that a source application component needs to be
2444 able to communicate with a target software component to consume the services of the target. [ConnectTo](#) is a
2445 component interdependency semantic in the most general sense and does not try imply how the communication
2446 between the source and target components is physically realized.

2447

2448 Application component intercommunication typically has conventions regarding which component(s) initiate the
2449 communication. Connection initiation semantics are specified in [tosca.capabilities.Endpoint](#). Endpoints at each
2450 end of the [tosca.relationships.ConnectsTo](#) must indicate identical connection initiation semantics.

2451

2452 The following sections describe the normative connection initiation semantics for the
2453 [tosca.relationships.ConnectsTo](#) Relationship Type.

2454 **F.3.1.1 Source to Target**

2455 The Source to Target communication initiation semantic is the most common case where the source component
2456 initiates communication with the target component in order to fulfill an instance of the

2457 tosca.relationships.ConnectsTo relationship. The typical case is a “client” component connecting to a “server”
2458 component where the client initiates a stream oriented connection to a pre-defined transport specific port or set of
2459 ports.

2460

2461 It is the responsibility of the TOSCA implementation to ensure the source component has a suitable network path
2462 to the target component and that the ports specified in the respective tosca.capabilities.Endpoint are not blocked.
2463 The TOSCA implementation may only represent state of the tosca.relationships.ConnectsTo relationship as
2464 fulfilled after the actual network communication is enabled and the source and target components are in their
2465 operational states.

2466

2467 Note that the connection initiation semantic only impacts the fulfillment of the actual connectivity and does not
2468 impact the node traversal order implied by the tosca.relationships.ConnectsTo Relationship Type.

2469 **F.3.1.2 Target to Source**

2470 The Target to Source communication initiation semantic is a less common case where the target component
2471 initiates communication with the source comment in order to fulfill an instance of the
2472 tosca.relationships.ConnectsTo relationship. This “reverse” connection initiation direction is typically required due
2473 to some technical requirements of the components or protocols involved, such as the requirement that SSH must
2474 only be initiated from target component in order to fulfill the services required by the source component.

2475

2476 It is the responsibility of the TOSCA implementation to ensure the source component has a suitable network path
2477 to the target component and that the ports specified in the respective tosca.capabilities.Endpoint are not blocked.
2478 The TOSCA implementation may only represent state of the tosca.relationships.ConnectsTo relationship as
2479 fulfilled after the actual network communication is enabled and the source and target components are in their
2480 operational states.

2481

2482 Note that the connection initiation semantic only impacts the fulfillment of the actual connectivity and does not
2483 impact the node traversal order implied by the tosca.relationships.ConnectsTo Relationship Type.

2484

2485 **F.3.1.3 Peer-to-Peer**

2486 The Peer-to-Peer communication initiation semantic allows any member of a group to initiate communication with
2487 any other member of the same group at any time. This semantic typically appears in clustering and distributed
2488 services where there is redundancy of components or services.

2489

2490 It is the responsibility of the TOSCA implementation to ensure the source component has a suitable network path
2491 between all the member component instances and that the ports specified in the respective
2492 tosca.capabilities.Endpoint are not blocked, and the appropriate multicast communication, if necessary, enabled.
2493 The TOSCA implementation may only represent state of the tosca.relationships.ConnectsTo relationship as
2494 fulfilled after the actual network communication is enabled such that at least one member component of the group
2495 may reach any other member component of the group.

2496

2497 Endpoints specifying the Peer-to-Peer initiation semantic need not be related with a
2498 tosca.relationships.ConnectsTo relationship for the common case where the same set of component instances
2499 must communicate with each other.

2500

2501 Note that the connection initiation semantic only impacts the fulfillment of the actual connectivity and does not
2502 impact the node traversal order implied by the tosca.relationships.ConnectsTo Relationship Type.

2503 **F.3.2 Specifying layer 4 ports**

2504 TOSCA Service Templates must express enough details about application component intercommunication to
2505 enable TOSCA implementations to fulfill these communication semantics in the network infrastructure. TOSCA
2506 currently focuses on TCP/IP as this is the most pervasive in today's cloud infrastructures. The layer 4 ports
2507 required for application component intercommunication are specified in `tosca.capabilities.Endpoint`. The union of
2508 the port specifications of both the source and target `tosca.capabilities.Endpoint` which are part of the
2509 `tosca.relationships.ConnectsTo` Relationship Template are interpreted as the effective set of ports which must be
2510 allowed in the network communication.

2511

2512 The meaning of Source and Target port(s) corresponds to the direction of the respective
2513 `tosca.relationships.ConnectsTo`.

2514 **F.4 Network provisioning**

2515 **F.4.1 Declarative network provisioning**

2516 TOSCA orchestrators are responsible for the provisioning of the network connectivity for declarative TOCSA
2517 Service Templates (Declarative TOCSA Service Templates don't contain explicit plans). This means that the
2518 TOSCA orchestrator must be able to infer a suitable logical connectivity model from the Service Template and
2519 then decide how to provision the logical connectivity, referred to as "fulfillment", on the available underlying
2520 infrastructure. In order to enable fulfillment, sufficient technical details still must be specified, such as the required
2521 protocols, ports and QOS information. TOSCA connectivity types, such as `tosca.capabilities.Endpoint`, provide
2522 well defined means to express these details.

2523 **F.4.2 Implicit network fulfillment**

2524 TOSCA Service Templates are by default network agnostic. TOSCA's application centric approach only requires
2525 that a TOSCA Service Template contain enough information for a TOSCA orchestrator to infer suitable network
2526 connectivity to meet the needs of the application components. Thus Service Template designers are not required
2527 to be aware of or provide specific requirements for underlying networks. This approach yields the most portable
2528 Service Templates, allowing them to be deployed into any infrastructure which can provide the necessary
2529 component interconnectivity.

2530 **F.4.3 Controlling network fulfillment**

2531 TOSCA provides mechanisms for providing control over network fulfillment.

2532 This mechanism allows the application network designer to express in service template or network template how
2533 the networks should be provisioned.

2534

2535 For the use cases described below let's assume we have a typical 3-tier application which is consisting of FE
2536 (frontend), BE (backend) and DB (database) tiers. The simple application topology diagram can be shown below:

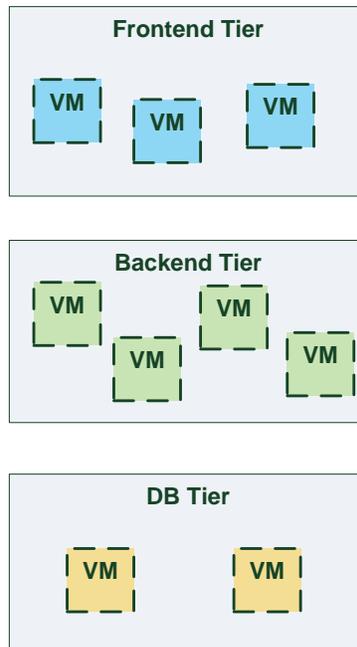


Figure-5: Typical 3-Tier Network

2537

2538

2539 **F.4.3.1 Use case: OAM Network**

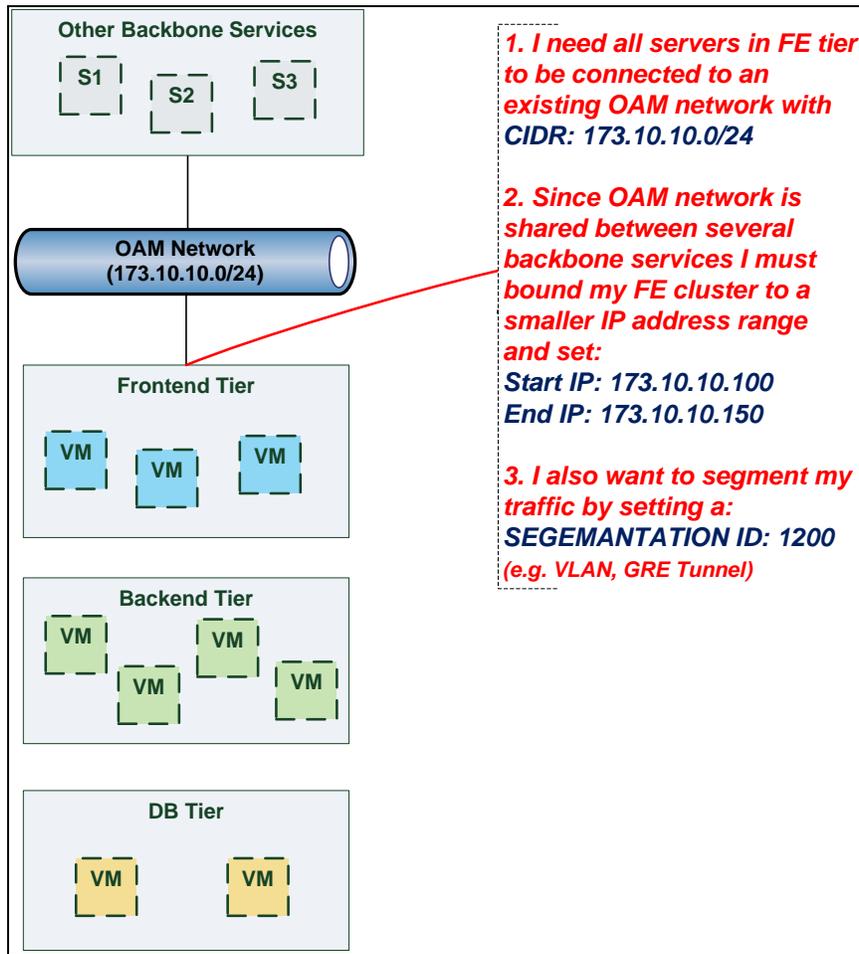
2540 When deploying an application in service provider's on-premise cloud, it's very common that one or more of the
 2541 application's services should be accessible from an ad-hoc OAM (Operations, Administration and Management)
 2542 network which exists in the service provider backbone.

2543

2544 As an application network designer, I'd like to express in my TOSCA network template (which corresponds to my
 2545 TOSCA service template) the network CIDR block, start ip, end ip and segmentation ID (e.g. VLAN id).

2546 The diagram below depicts a typical 3-tiers application with specific networking requirements for its FE tier server
 2547 cluster:

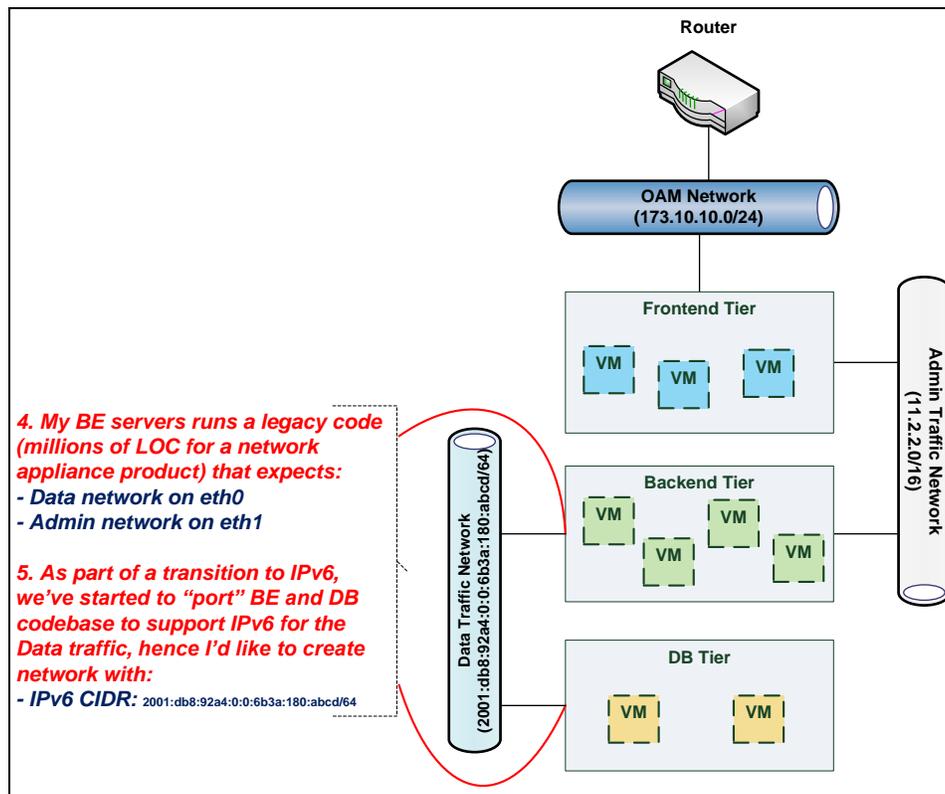
2548



2549

2550 **F.4.3.2 Use case: Data Traffic network**

2551 The diagram below defines a set of networking requirements for the backend and DB tiers of the 3-tier app
2552 mentioned above.



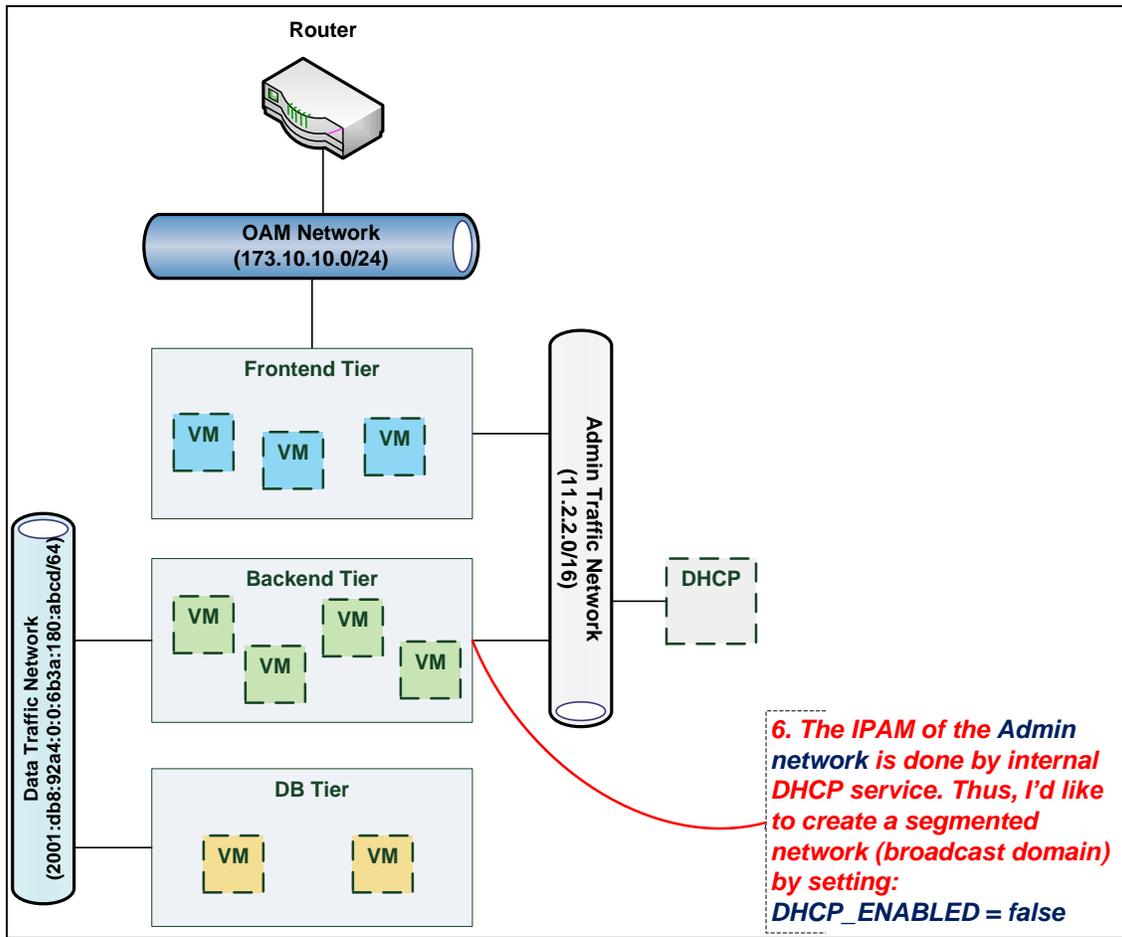
2553

2554 F.4.3.3 Use case: Bring my own DHCP

2555 The same 3-tier app requires for its admin traffic network to manage the IP allocation by its own DHCP which runs
 2556 autonomously as part of application domain.

2557

2558 For this purpose, the app network designer would like to express in TOSCA that the underlying provisioned
 2559 network will be set with DHCP_ENABLED=false. See this illustrated in the figure below:



2560

2561 F.5 Network Types

2562 F.5.1 toska.nodes.network.Network

2563 The TOSCA Network node represents a simple, logical network service.

Shorthand Name	Network
Type Qualified Name	tosca:Network
Type URI	tosca.nodes.network.Network

2564 F.5.1.1 Properties

Name	Required	Type	Constraints	Description
ip_version	no	integer	valid_values: [4, 6] default: 4	The IP version of the requested network
cidr	no	string	None	The cidr block of the requested network
start_ip	no	string	None	The IP address to be used as the 1 st one in a pool of addresses derived from the cidr block full IP range
end_ip	no	string	None	The IP address to be used as the last one in a pool of addresses derived from the cidr block full IP range
gateway_ip	no	string	None	The gateway IP address.

Name	Required	Type	Constraints	Description
network_name	no	string	None	<p>An Identifier that represents an existing Network instance in the underlying cloud infrastructure – OR – be used as the name of the new created network.</p> <ul style="list-style-type: none"> • If network_name is provided along with network_id they will be used to uniquely identify an existing network and not creating a new one, means all other possible properties are not allowed. • network_name should be more convenient for using. But in case that network name uniqueness is not guaranteed then one should provide a network_id as well.
network_id	no	string	None	<p>An Identifier that represents an existing Network instance in the underlying cloud infrastructure.</p> <p>This property is mutually exclusive with all other properties except network_name.</p> <ul style="list-style-type: none"> • Appearance of network_id in network template instructs the Tosca container to use an existing network instead of creating a new one. • network_name should be more convenient for using. But in case that network name uniqueness is not guaranteed then one should add a network_id as well. • network_name and network_id can be still used together to achieve both uniqueness and convenient.
segmentation_id	no	string	None	<p>A segmentation identifier in the underlying cloud infrastructure (e.g., VLAN id, GRE tunnel id). If the segmentation_id is specified, the network_type or physical_network properties should be provided as well.</p>
network_type	no	string	None	<p>Optionally, specifies the nature of the physical network in the underlying cloud infrastructure. Examples are flat, vlan, gre or vxlan. For flat and vlan types, physical_network should be provided too.</p>
physical_network	no	string	None	<p>Optionally, identifies the physical network on top of which the network is implemented, e.g. physnet1. This property is required if network_type is flat or vlan.</p>
dhcp_enabled	no	boolean	default: true	<p>Indicates the TOSCA container to create a virtual network instance with or without a DHCP service.</p>

2565 **F.5.1.2 Attributes**

Name	Required	Type	Constraints	Description
segmentation_id	no	string	None	<p>The actual <i>segmentation_id</i> that is been assigned to the network by the underlying cloud infrastructure.</p>

2566 **F.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

```

2567 **F.5.1.4 Additional Requirements**

- 2568 • None

2569 **F.5.2 `tosca.nodes.network.Port`**

2570 The TOSCA **Port** node represents a logical entity that associates between Compute and Network normative
 2571 types.

2572 The Port node type effectively represents a single virtual NIC on the Compute node instance.

Shorthand Name	Port
Type Qualified Name	tosca:Port
Type URI	tosca.nodes.network.Port

2573

F.5.2.1 Properties

Name	Required	Type	Constraints	Description
ip_address	no	string	None	Allow the user to set a fixed IP address. Note that this address is a request to the provider which they will attempt to fulfil but may not be able to dependent on the network the port is associated with.
order	no	integer	greater_or_equal: 0 default: 0	The order of the NIC on the compute instance (e.g. eth2). Note: when binding more than one port to a single compute (aka multi vNICs) and ordering is desired, it is <i>*mandatory*</i> that all ports will be set with an order value and. The <i>order</i> values must represent a positive, arithmetic progression that starts with 0 (e.g. 0, 1, 2, ..., n).
is_default	no	boolean	default: false	Set is_default =true to apply a default gateway route on the running compute instance to the associated network gateway. Only one port that is associated to single compute node can set as default=true.
ip_range_start	no	string	None	Defines the starting IP of a range to be allocated for the compute instances that are associated by this Port. Without setting this property the IP allocation is done from the entire CIDR block of the network.
ip_range_end	no	string	None	Defines the ending IP of a range to be allocated for the compute instances that are associated by this Port. Without setting this property the IP allocation is done from the entire CIDR block of the network.

2574

F.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.

2575

F.5.2.3 Definition

```

tosca.nodes.network.Port:
  derived_from: toasca.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

```

2576 **F.5.2.4 Additional Requirements**

- 2577
 - None

2578 **F.5.3 [tosca.capabilities.network.Linkable](#)**

2579 A node type that includes the Linkable capability indicates that it can be pointed by
 2580 [tosca.relationships.network.LinksTo](#) relationship type.

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

2581 **F.5.3.1 Properties**

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

2582 **F.5.3.2 Definition**

```

tosca.capabilities.network.Linkable:
  derived_from: tosca.capabilities.Node

```

2583 **F.5.4 [tosca.relationships.network.LinksTo](#)**

2584 This relationship type represents an association relationship between Port and Network node types.

Shorthand Name	LinksTo
Type Qualified Name	tosca:LinksTo
Type URI	tosca.relationships.network.LinksTo

2585 **F.5.4.1 Definition**

```
tosca.relationships.network.LinksTo:
  derived_from: tosca.relationships.DependsOn
  valid_target_types: [ tosca.capabilities.network.Linkable ]
```

2586 **F.5.5 [tosca.relationships.network.BindsTo](#)**

2587 This type represents a network association relationship between Port and Compute node types.

Shorthand Name	network.BindsTo
Type Qualified Name	tosca:BindsTo
Type URI	tosca.relationships.network.BindsTo

2588 **F.5.5.1 Definition**

```
tosca.relationships.network.BindsTo:
  derived_from: tosca.relationships.DependsOn
  valid_target_types: [ tosca.capabilities.network.Bindable ]
```

2589 **F.6 Network modeling approaches**

2590 **F.6.1 Option 1: Specifying a network outside the application’s Service Template**

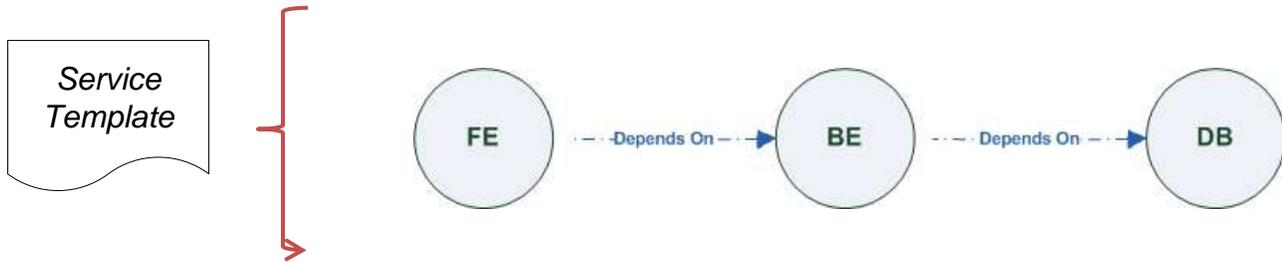
2591 This approach allows someone who understands the application’s networking requirements, mapping the details
2592 of the underlying network to the appropriate node templates in the application.

2593
2594 The motivation for this approach is providing the application network designer a fine-grained control on how
2595 networks are provisioned and stitched to its application by the TOSCA orchestrator and underlying cloud
2596 infrastructure while still preserving the portability of his service template. Preserving the portability means here not
2597 doing any modification in service template but just “plug-in” the desired network modeling. The network modeling
2598 can reside in the same service template file but the best practice should be placing it in a separated self-
2599 contained network template file.

2600
2601 This “pluggable” network template approach introduces a new normative node type called Port, capability called
2602 [tosca.capabilities.network.Linkable](#) and relationship type called [tosca.relationships.network.LinksTo](#).
2603 The idea of the Port is to elegantly associate the desired compute nodes with the desired network nodes while not
2604 “touching” the compute itself.

2605
2606 The following diagram series demonstrate the plug-ability strength of this approach.

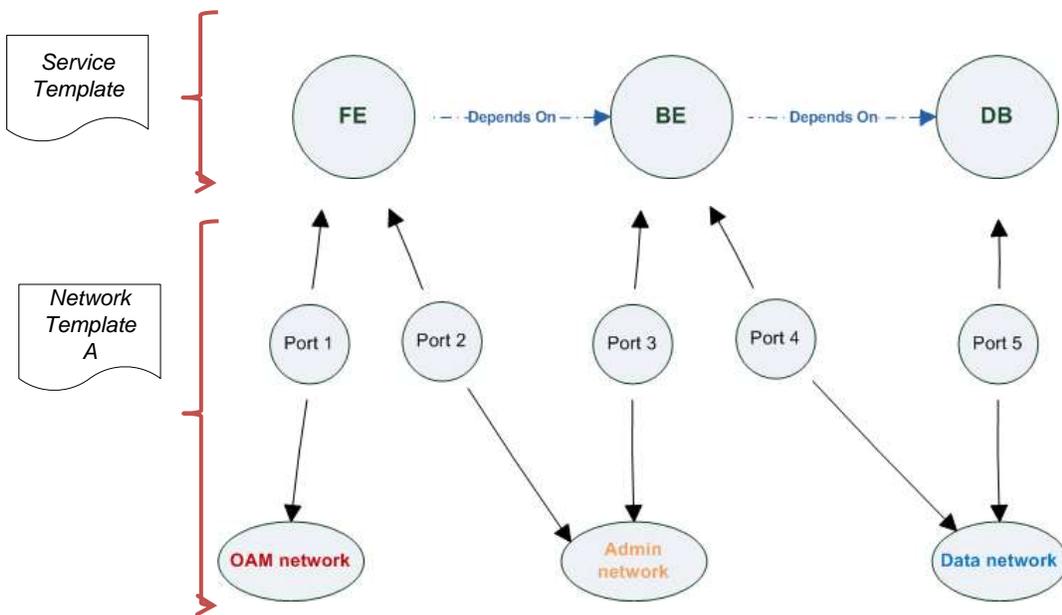
2607 Let's assume an application designer has modeled a service template as shown in Figure 1 that describes the
 2608 application topology nodes (compute, storage, software components, etc.) with their relationships. The
 2609 designer ideally wants to preserve this service template and use it in any cloud provider environment without
 2610 any change.



2611
 2612 *Figure-6: Generic Service Template*

2613 When the application designer comes to consider its application networking requirement they typically call the
 2614 network architect/designer from their company (who has the correct expertise).

2615 The network designer, after understanding the application connectivity requirements and optionally the target
 2616 cloud provider environment, is able to model the network template and plug it to the service template as shown
 2617 in Figure 2:



2618
 2619 *Figure-7: Service template with network template A*

2620 When there's a new target cloud environment to run the application on, the network designer is simply creates
 2621 a new network template B that corresponds to the new environmental conditions and provide it to the
 2622 application designer which packs it into the application CSAR.

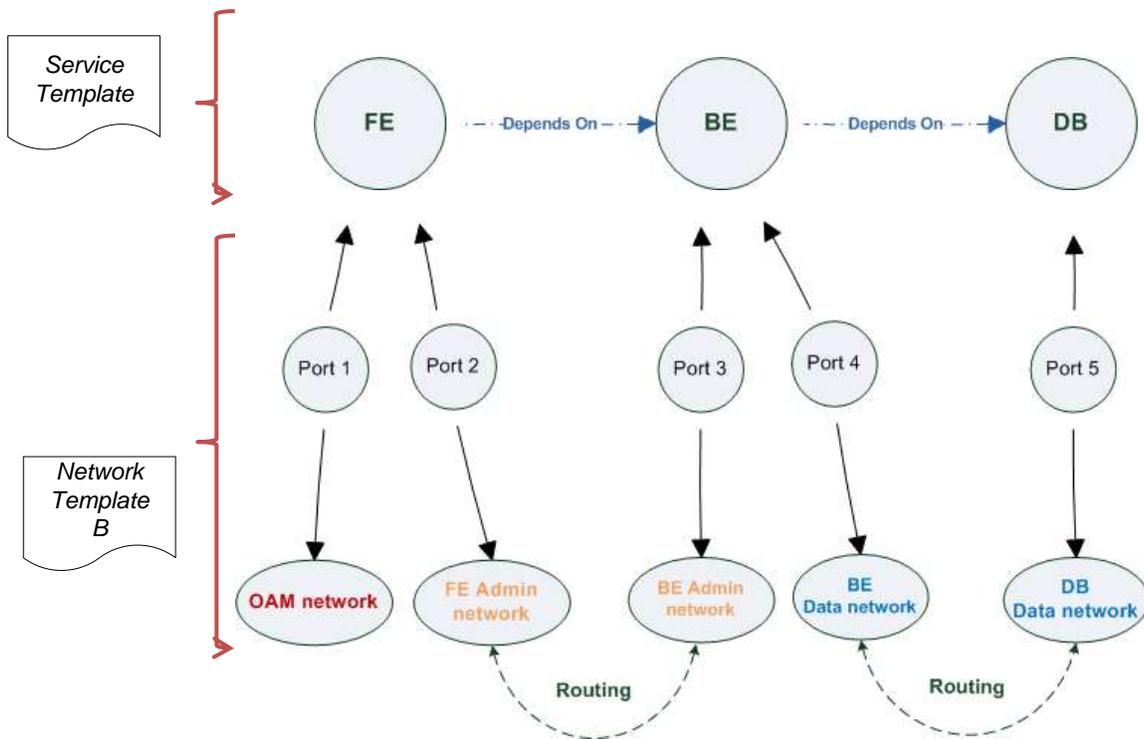


Figure-8: Service template with network template B

2623

2624

2625

The node templates for these three networks would be defined as follows:

```

node_templates:
  frontend:
    type: toasca.nodes.Compute
    properties: # omitted for brevity

  backend:
    type: toasca.nodes.Compute
    properties: # omitted for brevity

  database:
    type: toasca.nodes.Compute
    properties: # omitted for brevity

  oam_network:
    type: toasca.nodes.network.Network
    properties: # omitted for brevity

  admin_network:
    type: toasca.nodes.network.Network
    properties: # omitted for brevity

  data_network:
  
```

```
type: 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
```

2626
2627
2628
2629
2630

F.6.2 Option 2: Specifying network requirements within the application's Service Template

This approach allows the Service Template designer to map an endpoint to a logical network.

The use case shown below examines a way to express in the TOSCA YAML service template a typical 3-tier application with their required networking modeling:

```
node_templates:
  frontend:
    type: toasca.nodes.Compute
    properties: # omitted for brevity
    requirements:
      - network_oam: oam_network
      - network_admin: admin_network
  backend:
    type: toasca.nodes.Compute
    properties: # omitted for brevity
    requirements:
      - network_admin: admin_network
      - network_data: data_network

  database:
    type: toasca.nodes.Compute
    properties: # omitted for brevity
    requirements:
      - network_data: data_network

  oam_network:
    type: toasca.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: toasca.nodes.network.Network
    properties:
      ip_version: { get_input: admin_network_ip_version }
      dhcp_enabled: { get_input: admin_network_dhcp_enabled }

  data_network:
    type: toasca.nodes.network.Network
```

```
properties:  
  ip_version: { get_input: data_network_ip_version }  
  cidr: { get_input: data_network_cidr }
```

2631

Appendix G. Component Modeling Use Cases

G.1.1 Use Case: Exploring the HostedOn relationship using WebApplication and WebServer

This use case examines the ways TOSCA YAML can be used to express a simple hosting relationship (i.e., **HostedOn**) using the normative TOSCA **WebServer** and **WebApplication** node types defined in this specification.

G.1.1.1 WebServer declares its “host” capability

For convenience, relevant parts of the normative TOSCA Node Type for **WebServer** are shown below:

```
tosca.nodes.WebServer
  derived_from: SoftwareComponent
  capabilities:
    ...
  host:
    type: tosca.capabilities.Container
    valid_source_types: [ tosca.nodes.WebApplication ]
```

As can be seen, the **WebServer** Node Type declares its capability to “contain” (i.e., host) other nodes using the symbolic name “**host**” and providing the Capability Type **tosca.capabilities.Container**. It should be noted that the symbolic name of “**host**” is not a reserved word, but one assigned by the type designer that implies at or betokens the associated capability. The **Container** capability definition also includes a required list of valid Node Types that can be contained by this, the **WebServer**, Node Type. This list is declared using the keyname of **valid_source_types** and in this case it includes only allowed type **WebApplication**.

G.1.1.2 WebApplication declares its “host” requirement

The **WebApplication** node type needs to be able to describe the type of capability a target node would have to provide in order to “host” it. The normative TOSCA capability type **tosca.capabilities.Container** is used to describe all normative TOSCA hosting (i.e., container-containee pattern) relationships. As can be seen below, the **WebApplication** accomplishes this by declaring a requirement with the symbolic name “**host**” with the **capability** keyname set to **tosca.capabilities.Container**.

Again, for convenience, the relevant parts of the normative **WebApplication** Node Type are shown below:

```
tosca.nodes.WebApplication:
  derived_from: tosca.nodes.Root
  requirements:
    - host:
        capability: tosca.capabilities.Container
        node: tosca.nodes.WebServer
        relationship: tosca.relationships.HostedOn
```

G.1.1.2.1 Notes

- 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

2655 HostedOn relationship could still be established between WebApplicaton and WebServer in a TOSCA
2656 Service Template regardless of the symbolic name assigned to either the requirement or capability
2657 declaration.

2658 **G.1.2 Use Case: Establishing a ConnectsTo relationship to WebServer**

2659 This use case examines the ways TOSCA YAML can be used to express a simple connection relationship (i.e.,
2660 [ConnectsTo](#)) between some service derived from the [SoftwareComponent](#) Node Type, to the normative
2661 [WebServer](#) node type defined in this specification.

2662 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
```

2663 Since the normative [WebServer](#) Node Type only declares one capability of type [tosca.capabilities.Endpoint](#)
2664 (or [Endpoint](#), its alias in TOSCA) using the symbolic name [data_endpoint](#), the [my_web_service](#) node template
2665 does not need to declare that symbolic name on its requirement declaration. If however, the [my_web_server](#)
2666 node was based upon some other node type that declared more than one capability of type [Endpoint](#), then the
2667 **capability** keyname could be used to supply the desired symbolic name if necessary.

2668 **G.1.2.1 Best practice**

2669 It should be noted that the best practice for designing Node Types in TOSCA should not export two capabilities
2670 of the same type if they truly offer different functionality (i.e., different capabilities) which should be
2671 distinguished using different Capability Type definitions.

2672 **G.1.3 Use Case: Attaching (local) BlockStorage to a Compute node**

2673 This use case examines the ways TOSCA YAML can be used to express a simple AttachesTo relationship
2674 between a Compute node and a locally attached BlockStorage node.

2675 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
```

2676 **G.1.4 Use Case: Reusing a BlockStorage Relationship using Relationship Type or**
2677 **Relationship Template**

2678 This builds upon the previous use case (G.1.3) to examine how a template author could attach multiple Compute
2679 nodes (templates) to the same BlockStorage node (template), but with slightly different property values for the
2680 AttachesTo relationship.

2681
2682 Specifically, several notation options are shown (in this use case) that achieve the same desired result.

2683 **G.1.4.1 Simple Profile Rationale**

2684 Referencing an explicitly declared Relationship Template is a convenience of the Simple Profile that allows
2685 template authors an entity to set, constrain or override the properties and operations as defined in its declared
2686 (Relationship) Type much as allowed now for Node Templates. It is especially useful when a complex
2687 Relationship Type (with many configurable properties or operations) has several logical occurrences in the same
2688 Service (Topology) Template; allowing the author to avoid configuring these same properties and operations in
2689 multiple Node Templates.

2690 **G.1.4.2 Notation Style #1: Augment AttachesTo Relationship Type directly in each Node**
2691 **Template**

2692 This notation extends the methodology used for establishing a HostedOn relationship, but allowing template
2693 author to supply (dynamic) configuration and/or override of properties and operations.

2694

2695

2696

2697

Note: This option will remain valid for Simple Profile regardless of other (following) notation (or aliasing) options being discussed or adopted.

```
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
```

2698

2699
2700

G.1.4.3 Notation Style #2: Use the 'template' keyword on the Node Templates to specify which named Relationship Template to use

2701
2702
2703
2704

This option shows how to explicitly declare different named Relationship Templates within the Service Template as part of a `relationship_templates` section (which have different property values) and can be referenced by different Compute typed Node Templates.

```
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

  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
```

2705

2706

G.1.4.4 Notation Style #3: Using the “copy” keyname to define a similar Relationship Template

2707

2708

How does TOSCA make it easier to create a new relationship template that is mostly the same as one that exists without manually copying all the same information? TOSCA provides the **copy** keyname as a convenient way to copy an existing template definition into a new template definition as a starting point or basis for describing a new definition and avoid manual copy. The end results are cleaner TOSCA Service Templates that allows the description of only the changes (or deltas) between similar templates.

2709

2710

2711

2712

2713

The example below shows that the Relationship Template named **storage_attachesto_1** provides some

2714

overrides (conceptually a large set of overrides) on its Type which the Relationship Template named

2715

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
```

2716

Appendix H. Complete Application Modeling Use Cases

2717

H.1 Use cases

2718

H.1.1 Overview

Use Case		
Name	Description	Service Template link (Entry Definition)
Compute: Hosting a Virtual Machine (VM) on a single Compute instance	Introduces a TOSCA Compute node which is used to stand up a single instance of a Virtual Machine (VM) image.	TODO
BlockStorage-1: Attaching Block Storage to a single Compute instance	Demonstrates how to attach a TOSCA BlockStorage node to a Compute node using the normative AttachesTo relationship.	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/e/tosca_blockstorage_with_attachment.yaml#L19
BlockStorage-2: Attaching Block Storage using a custom Relationship Type	Demonstrates how to attach a TOSCA BlockStorage node to a Compute node using a custom RelationshipType that derives from the normative AttachesTo relationship.	TODO
BlockStorage-3: Using a Relationship Template of type AttachesTo	Demonstrates how to attach a TOSCA BlockStorage node to a Compute node using a TOSCA Relationship Template that is based upon the normative AttachesTo Relationship Type.	TODO
BlockStorage-4: Single Block Storage shared by 2-Tier Application with custom AttachesTo Type and implied relationships	This use case shows 2 compute instances (2 tiers) with one BlockStorage node, and also uses a custom AttachesTo Relationship that provides a default mount point (i.e., location) which the 1 st tier uses, but the 2 nd tier provides a different mount point.	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/e/tosca_blockstorage_with_attachment_notation1.yaml
BlockStorage-5: Single Block Storage shared by 2-Tier Application with custom AttachesTo Type and explicit Relationship Templates	This use case is like the previous BlockStorage-4 use case, but also creates two relationship templates (one for each tier) each of which provide a different mount point (i.e., location) which overrides the default location defined in the custom Relationship Type.	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/e/tosca_blockstorage_with_attachment_notation2.yaml
BlockStorage-6: Multiple Block Storage attached to different Servers	This use case demonstrates how two different TOSCA BlockStorage nodes can be attached to two different Compute nodes (i.e., servers) each using the normative AttachesTo relationship.	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/e/tosca_multiple_blockstorage_with_attachment.yaml
Object Storage 1: Creating an Object Storage service	Introduces the TOSCA ObjectStorage node type and shows how it can be instantiated.	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/e/tosca_single_object_store.yaml
Network-1: Server bound to a new network	Introduces the TOSCA Network and Port nodes used for modeling logical networks using the LinksTo and BindsTo Relationship Types. In this use case, the template is invoked without an existing network_name as an input property so a new network is created using the properties declared in the Network node.	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/rk/tosca_one_server_one_network.yaml
Network-2: Server bound to an existing network	Shows how to use a network_name as an input parameter to the template to allow a server to be associated with (i.e. bound to) an existing Network .	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/rk/tosca_server_on_existing_network.yaml
Network-3: Two servers bound to a single network	This use case shows how two servers (Compute nodes) can be associated with the same Network node using two	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/rk/tosca_two_servers_on_existing_network.yaml

Use Case		
Name	Description	Service Template link (Entry Definition)
	logical network Ports .	rk/tosca_two_servers_one_network.yaml
Network-4: Server bound to three networks	This use case shows how three logical networks (Network nodes), each with its own IP address range, can be associated with the same server (Compute node).	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/rk/tosca_one_server_three_networks.yaml
WebServer-DBMS-1: WordPress + MySQL, single instance	Shows how to host a TOSCA WebServer with a TOSCA WebApplication, DBMS and Database Node Types along with their dependent HostedOn and ConnectsTo relationships.	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/single_instance_wordpress.yaml
WebServer-DBMS-2: WordPress + MySQL + Floating IPs, single instance	Shows the WordPress web application and MySQL database nodes hosted on a single server (instance) along with demonstrating how to create a network for the application with Floating IP addresses.	TODO
WebServer-DBMS-3: Nodejs with PayPal Sample App and MongoDB on separate instances	Instantiates a 2-tier application with Nodejs and its (PayPal sample) WebApplication on one tier which connects a MongoDB database (which stores its application data) using a ConnectsTo relationship.	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/nodejs_mongodb_two_instances.yaml
Multi-Tier-1: Elasticsearch, Logstash, Kibana (ELK)	Shows Elasticsearch, Logstash and Kibana (ELK) being used in a typical manner to collect, search and monitor/visualize data from a running application. This use case builds upon the previous Nodejs/MongoDB 2-tier application as the one being monitored. The collectd and rsyslog components are added to both the WebServer and Database tiers which work to collect data for Logstash. In addition to the application tiers, a 3 rd tier is introduced with Logstash to collect data from the application tiers. Finally a 4 th tier is added to search the Logstash data with Elasticsearch and visualize it using Kibana . Note: This use case also shows the convenience of using a single YAML macro (declared in the dsl_definitions section of the TOSCA Service Template) on multiple Compute nodes.	https://github.com/openstack/heat-translator/blob/master/translator/toscalib/tests/data/elk.yaml
Container-1: Containers using Docker single Compute instance (Containers only)	Minimalist TOSCA Service Template description of 2 Docker containers linked to each other. Specifically, one container runs wordpress and connects to second mysql database container both on a single server (i.e., Compute instance). The use case also demonstrates how TOSCA declares and references Docker images from the Docker Hub repository. Variation 1: Docker Container nodes (only) providing their Docker Requirements allowing platform (orchestrator) to select/provide the underlying Docker implementation (Capability).	TODO

2719 H.1.2 Compute: Hosting a Virtual Machine (VM) on a single instance

2720 H.1.2.1 Description

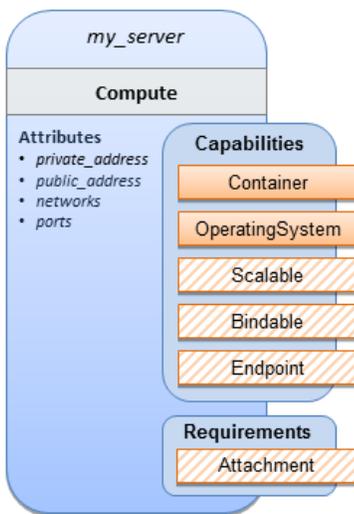
2721 This use case demonstrates how the TOSCA Simple Profile specification can be used to stand up a single
2722 instance of a Virtual Machine (VM) image using a normative TOSCA **Compute** node. The TOSCA Compute node
2723 is declarative in that the service template describes both the processor and host operating system platform
2724 characteristics (i.e., properties declared on the capability named “os”) that are desired by the template author.
2725 The cloud provider would attempt to fulfill these properties (to the best of its abilities) during orchestration.

2726 H.1.2.2 Features

2727 This use case introduces the following TOSCA Simple Profile features:

- 2728 • A node template that uses the normative TOSCA **Compute** Node Type along with showing an exemplary
2729 set of its properties being configured.
- 2730 • Use of the TOSCA Service Template **inputs** section to declare a configurable value the template user
2731 may supply at runtime. In this case, the “host” property named “num_cpus” (of type integer) is
2732 declared.
 - 2733 ○ Use of a property constraint to limit the allowed integer values for the “num_cpus” property to a
2734 specific list supplied in the property declaration.
- 2735 • Use of the TOSCA Service Template **outputs** section to declare a value the template user may request at
2736 runtime. In this case, the property named “instance_ip” is declared
 - 2737 ○ The “instance_ip” output property is programmatically retrieved from the **Compute** node’s
2738 “public_address” attribute using the TOSCA Service Template-level **get_attribute** function.

2739 H.1.2.3 Logical Diagram



2740

2741 H.1.2.4 Sample YAML

```
tosca_definitions_version: tosca_simple_yaml_1_0_0
```

```
description: >
```

```
    TOSCA simple profile that just defines a single compute instance. 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: 4 MB
          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] }

```

2742 **H.1.2.5 Notes**

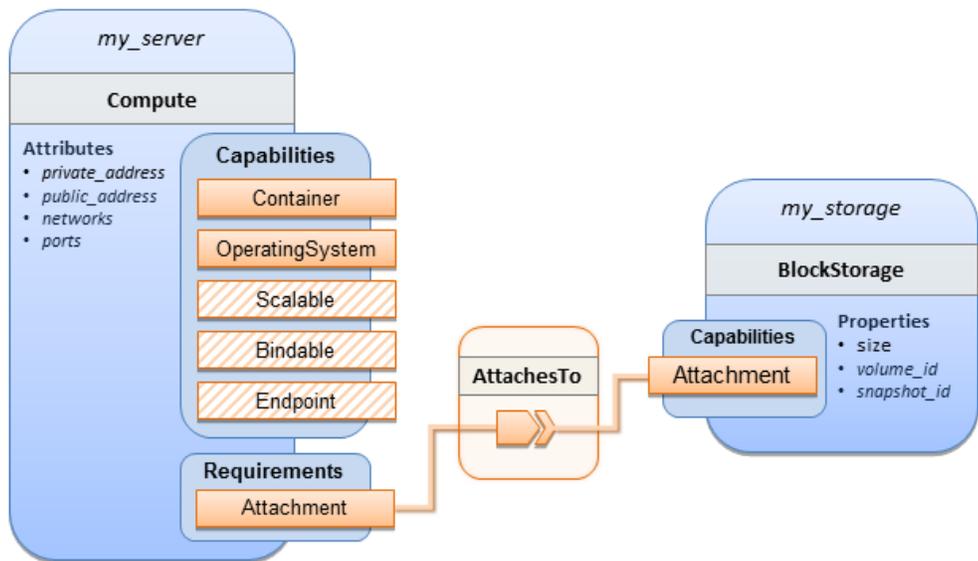
- 2743
 - This use case uses a versioned, Linux Ubuntu distribution on the Compute node.

2744 **H.1.3 Block Storage 1: Using the normative AttachesTo Relationship Type**

2745 **H.1.3.1 Description**

2746 This use case demonstrates how to attach a TOSCA **BlockStorage** node to a **Compute** node using the
 2747 normative **AttachesTo** relationship.

2748 **H.1.3.2 Logical Diagram**



2749

2750 **H.1.3.3 Sample YAML**

```

tosca_definitions_version: tosca_simple_yaml_1_0_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: 4096 kB
        os:
          properties:
            architecture: x86_64
            type: linux
            distribution: fedora
            version: 18.0
      requirements:
        - attachment:
            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] }

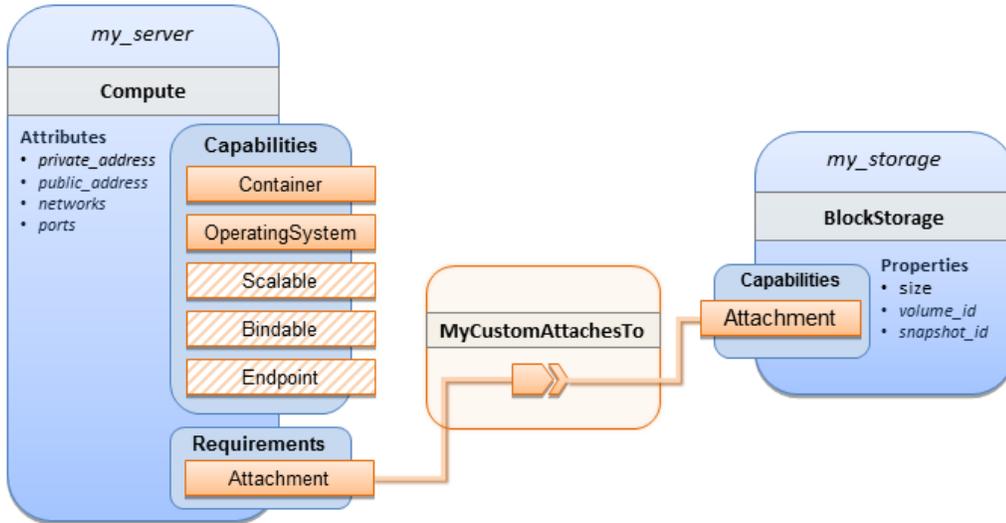
```

2751 H.1.4 Block Storage 2: Using a custom AttachesTo Relationship Type

2752 H.1.4.1 Description

2753 This use case demonstrates how to attach a TOSCA **BlockStorage** node to a **Compute** node using a custom
 2754 RelationshipType that derives from the normative **AttachesTo** relationship.

2755 **H.1.4.2 Logical Diagram**



2756

2757 **H.1.4.3 Sample YAML**

2758

```

tosca_definitions_version: toska_simple_yaml_1_0_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 MB
        os:
          properties:
            architecture: x86_64
            type: Linux
            distribution: Fedora
            version: 18.0
      requirements:
        - attachment:
            node: 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] }

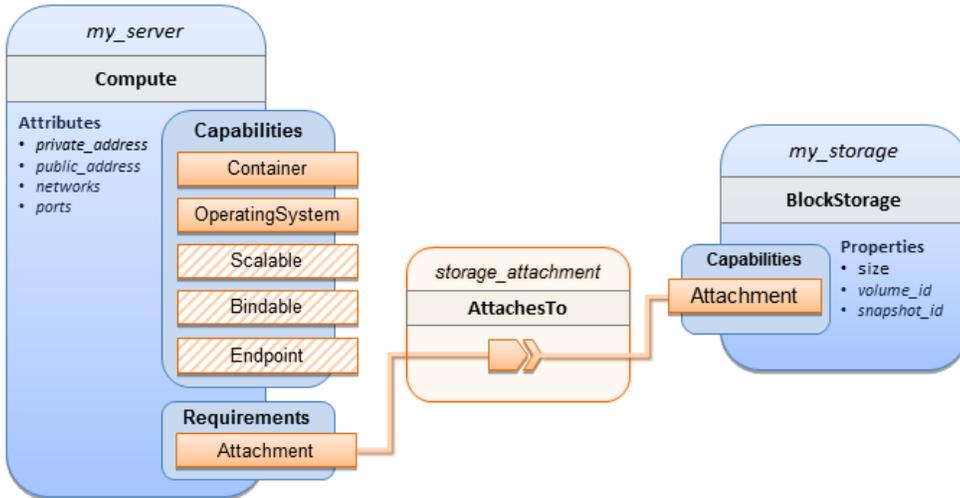
```

2759 **H.1.5 Block Storage 3: Using a Relationship Template of type AttachesTo**

2760 **H.1.5.1 Description**

2761 This use case demonstrates how to attach a TOSCA **BlockStorage** node to a **Compute** node using a TOSCA
2762 Relationship Template that is based upon the normative **AttachesTo** Relationship Type.

2763 **H.1.5.2 Logical Diagram**



2764

2765 **H.1.5.3 Sample YAML**

2766

```

tosca_definitions_version: tosca_simple_yaml_1_0_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 MB
        os:
          properties:
            architecture: x86_64
            type: Linux
            distribution: Fedora
            version: 18.0
      requirements:
        - local_storage:
            node: 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] }

```

2767
2768

H.1.6 Block Storage 4: Single Block Storage shared by 2-Tier Application with custom AttachesTo Type and implied relationships

2769

H.1.6.1 Description

2770
2771
2772

This use case shows 2 compute instances (2 tiers) with one BlockStorage node, and also uses a custom **AttachesTo** Relationship that provides a default mount point (i.e., **location**) which the 1st tier uses, but the 2nd tier provides a different mount point.

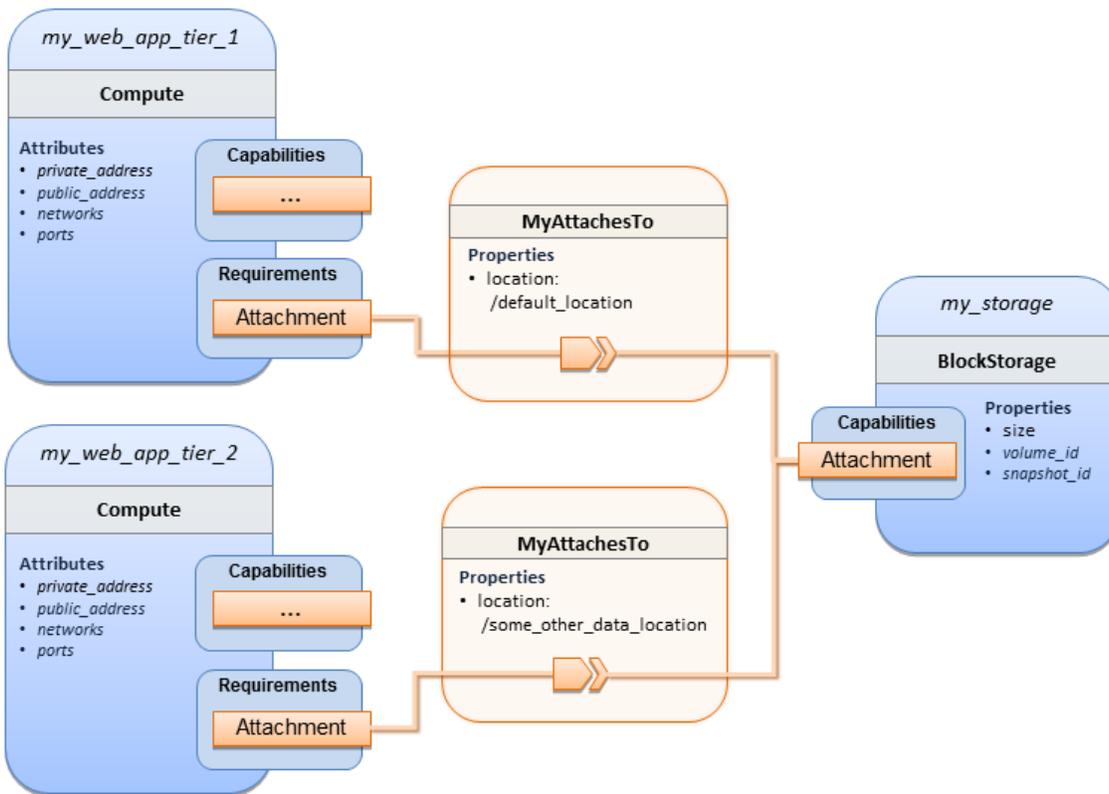
2773

2774
2775

Please note that this use case assumes both Compute nodes are accessing different directories within the shared, block storage node to avoid collisions.

2776

H.1.6.2 Logical Diagram



2777

2778

H.1.6.3 Sample YAML

```

tosca_definitions_version: tosca_simple_yaml_1_0_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:
          - attachment:
              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:
    - attachment:
        node: my_storage
        relationship:
          type: MyAttachesTo
          properties:
            location: /some_other_data_location

my_storage:
  type: toska.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] }

```

2779
2780

H.1.7 Block Storage 5: Single Block Storage shared by 2-Tier Application with custom AttachesTo Type and explicit Relationship Templates

2781

H.1.7.1 Description

2782
2783
2784

This use case is like the Notation1 use case, but also creates two relationship templates (one for each tier) each of which provide a different mount point (i.e., **location**) which overrides the default location defined in the custom Relationship Type.

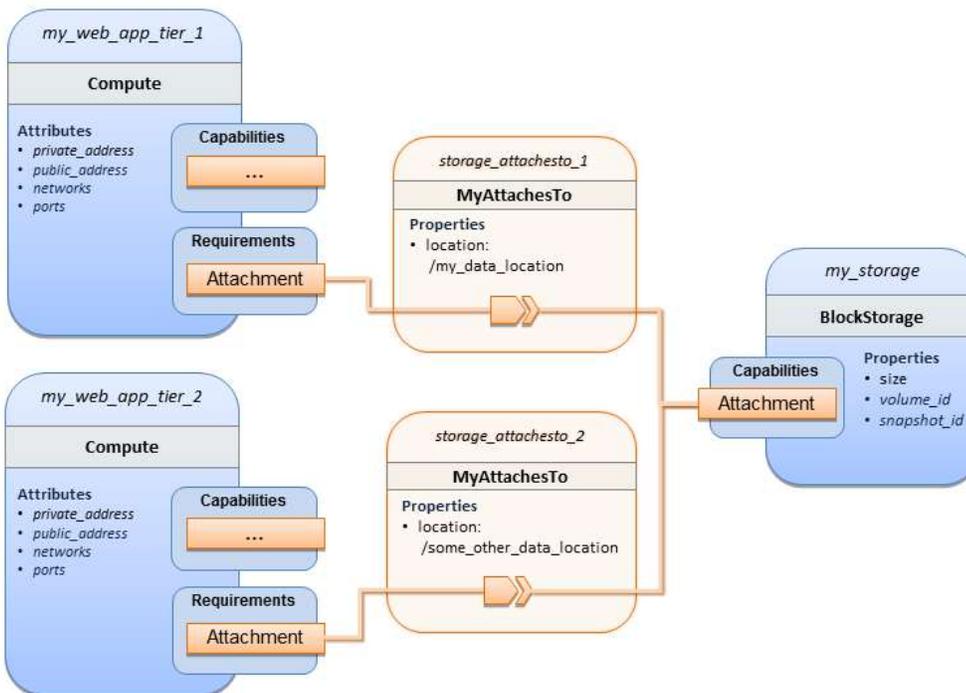
2785

2786
2787

Please note that this use case assumes both Compute nodes are accessing different directories within the shared, block storage node to avoid collisions.

2788

H.1.7.2 Logical Diagram



2789

2790

H.1.7.3 Sample YAML

```
tosca_definitions_version: tosca_simple_yaml_1_0_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 kB
    os:
      properties:
        architecture: x86_64
        type: Linux
        distribution: Fedora
        version: 18.0
  requirements:
    - attachment:
        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 kB
      os:
        properties:
          architecture: x86_64
          type: Linux
          distribution: Fedora
          version: 18.0
    requirements:
      - attachment:
          node: my_storage
          relationship: storage_attachesto_2

my_storage:
  type: toska.nodes.BlockStorage
  properties:
    size: { get_input: storage_size }
    snapshot_id: { get_input: storage_snapshot_id }

relationship_templates:
  storage_attachesto_1:
    type: MyAttachesTo
    properties:
      location: /my_data_location

  storage_attachesto_2:
    type: MyAttachesTo
    properties:
      location: /some_other_data_location

outputs:
  private_ip_1:
    description: The private IP address of the application's first tier.
    value: { get_attribute: [my_web_app_tier_1, private_address] }
  private_ip_2:
    description: The private IP address of the application's second tier.

```

```

value: { get_attribute: [my_web_app_tier_2, private_address] }
volume_id:
description: The volume id of the block storage instance.
value: { get_attribute: [my_storage, volume_id] }

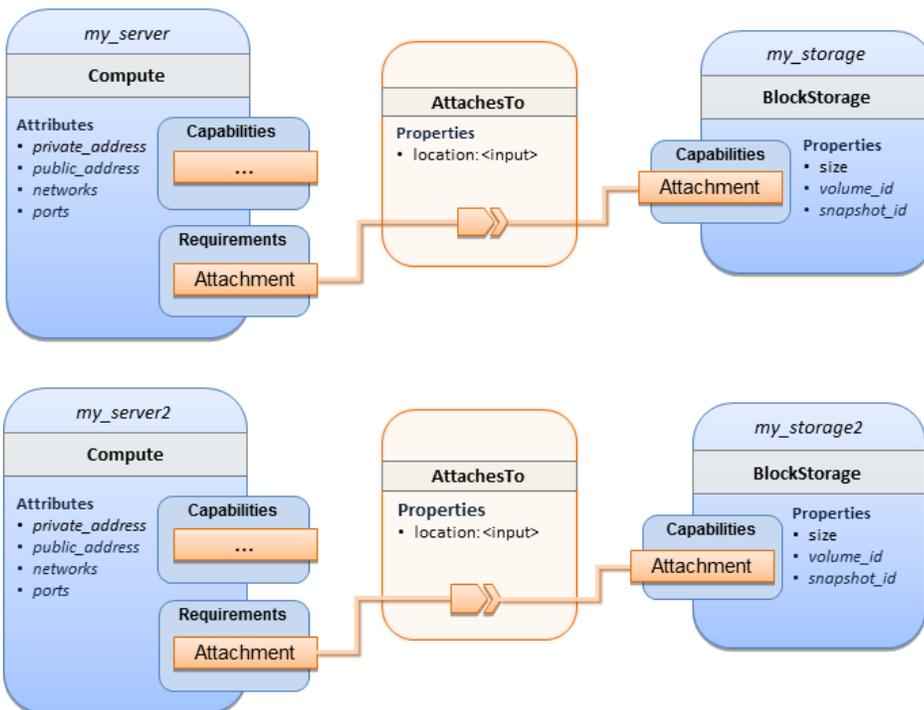
```

2791 **H.1.8 Block Storage 6: Multiple Block Storage attached to different Servers**

2792 **H.1.8.1 Description**

2793 This use case demonstrates how two different TOSCA **BlockStorage** nodes can be attached to two different
2794 **Compute** nodes (i.e., servers) each using the normative **AttachesTo** relationship.

2795 **H.1.8.2 Logical Diagram**



2796

2797 **H.1.8.3 Sample YAML**

```

tosca_definitions_version: tosca_simple_yaml_1_0_0

description: >
  TOSCA simple profile with 2 servers each with different attached block storage.

topology_template:
  inputs:
    cpus:
      type: integer
      description: Number of CPUs for the server.

```

```

constraints:
  - valid_values: [ 1, 2, 4, 8 ]
storage_size:
  type: scalar-unit.size
  default: 1 GB
  description: Size of the storage to be created.
storage_snapshot_id:
  type: string
  description: >
    Optional identifier for an existing snapshot to use when creating storage.
storage_location:
  type: string
  description: >
    Block storage mount point (filesystem path).

node_templates:
  my_server:
    type: tosca.nodes.Compute
    capabilities:
      host:
        properties:
          disk_size: 10 GB
          num_cpus: { get_input: cpus }
          mem_size: 4096 MB
      os:
        properties:
          architecture: x86_64
          type: Linux
          distribution: Fedora
          version: 18.0
    requirements:
      - attachment:
          node: my_storage
          relationship: AttachesTo
          properties:
            location: { get_input: storage_location }
  my_storage:
    type: tosca.nodes.BlockStorage
    properties:
      size: { get_input: storage_size }
      snapshot_id: { get_input: storage_snapshot_id }

```

```

my_server2:
  type: 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:
        - attachment:
            node: my_storage2
            relationship: AttachesTo
            properties:
              location: { get_input: storage_location }
my_storage2:
  type: toska.nodes.BlockStorage
  properties:
    size: { get_input: storage_size }
    snapshot_id: { get_input: storage_snapshot_id }

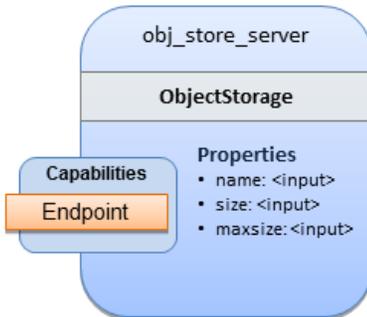
outputs:
  server_ip_1:
    description: The private IP address of the application's first server.
    value: { get_attribute: [my_server, private_address] }
  server_ip_2:
    description: The private IP address of the application's second server.
    value: { get_attribute: [my_server2, private_address] }
  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] }

```

2798 H.1.9 Object Storage 1: Creating an Object Storage service

2799 H.1.9.1 Description

2800 H.1.9.2 Logical Diagram



2801

2802 H.1.9.3 Sample YAML

```
tosca_definitions_version: tosca_simple_yaml_1_0_0

description: >
  Tosca template for creating an object storage service.

topology_template:
  inputs:
    objectstore_name:
      type: string

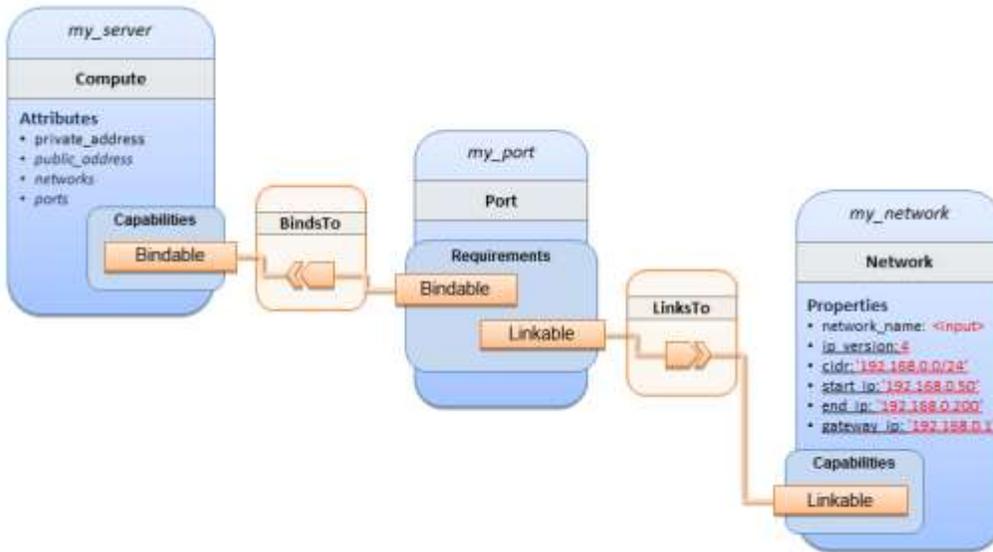
  node_templates:
    obj_store_server:
      type: tosca.nodes.ObjectStorage
      properties:
        name: { get_input: objectstore_name }
        size: 1024 kB
        maxsize: 1 GB
```

2803 H.1.10 Network 1: Server bound to a new network

2804 H.1.10.1 Description

2805 Introduces the TOSCA **Network** and **Port** nodes used for modeling logical networks using the **LinksTo** and **BindsTo**
2806 Relationship Types. In this use case, the template is invoked without an existing `network_name` as an input
2807 property so a new network is created using the properties declared in the Network node.

2808 **H.1.10.2 Logical Diagram**



2809

2810 **H.1.10.3 Sample YAML**

```
tosca_definitions_version: tosca_simple_yaml_1_0_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: 512 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'
```

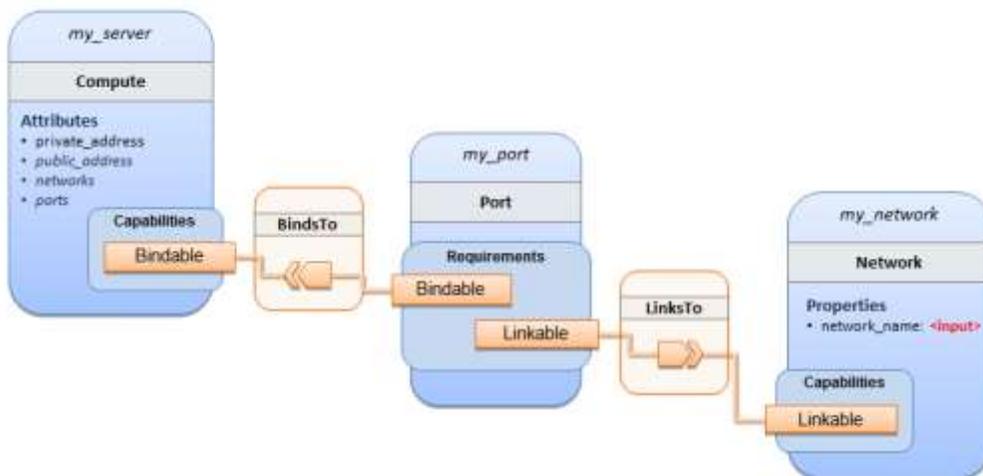
```
my_port:
  type: toska.nodes.network.Port
  requirements:
    - binding: my_server
    - link: my_network
```

2811 H.1.11 Network 2: Server bound to an existing network

2812 H.1.11.1 Description

2813 This use case shows how to use a `network_name` as an input parameter to the template to allow a server to be
2814 associated with an existing network.

2815 H.1.11.2 Logical Diagram



2816

2817 H.1.11.3 Sample YAML

```
tosca_definitions_version: toska_simple_yaml_1_0_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
      properties:
        disk_size: 10
        num_cpus: 1
        mem_size: 512
      capabilities:
        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

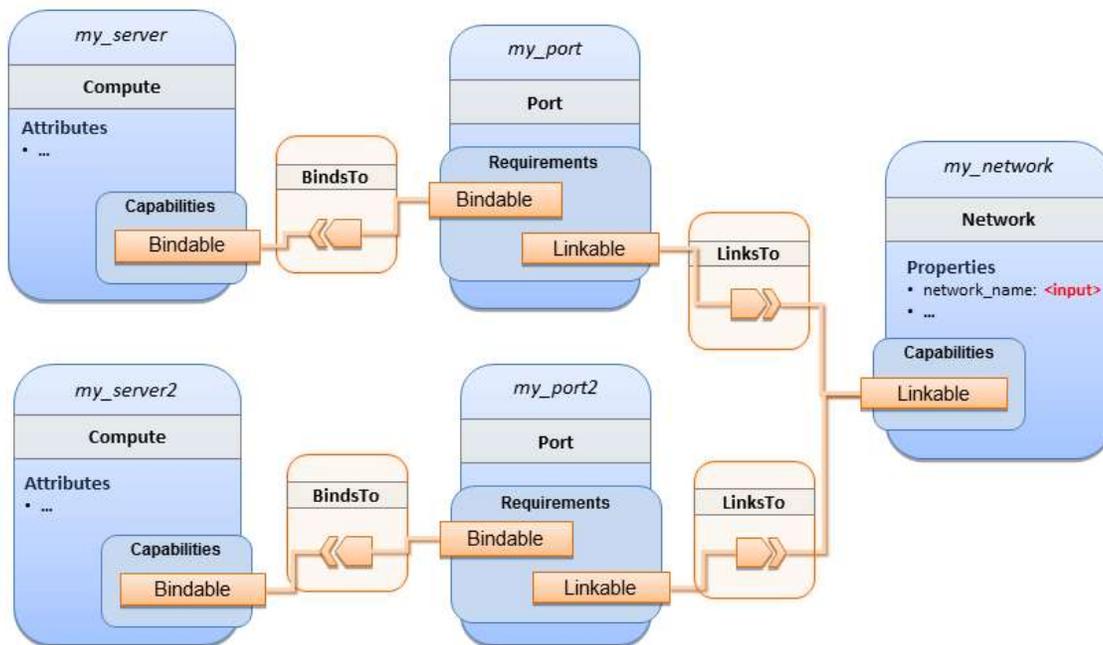
```

2818 H.1.12 Network 3: Two servers bound to a single network

2819 H.1.12.1 Description

2820 This use case shows how two servers (**Compute** nodes) can be bound to the same **Network** (node) using two
 2821 logical network **Ports**.

2822 **H.1.12.2 Logical Diagram**



2823

2824 **H.1.12.3 Sample YAML**

```
tosca_definitions_version: tosca_simple_yaml_1_0_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: toska.nodes.Compute
    capabilities:
      host:
        properties:
          disk_size: 10 GB
          num_cpus: 1
          mem_size: 512 MB
      os:
        properties:
          architecture: x86_64
          type: Linux
          distribution: CirrOS
          version: 0.3.2

  my_server2:
    type: toska.nodes.Compute
    capabilities:
      host:
        properties:
          disk_size: 10 GB
          num_cpus: 1
          mem_size: 512 MB
      os:
        properties:
          architecture: x86_64
          type: Linux
          distribution: CirrOS
          version: 0.3.2

  my_network:
    type: toska.nodes.network.Network
    properties:
      ip_version: 4
      cidr: { get_input: network_cidr }
      network_name: { get_input: network_name }
      start_ip: { get_input: network_start_ip }
      end_ip: { get_input: network_end_ip }

  my_port:
```

```

type: toska.nodes.network.Port
requirements:
  - binding: my_server
  - link: my_network

```

```

my_port2:
type: toska.nodes.network.Port
requirements:
  - binding: my_server2
  - link: my_network

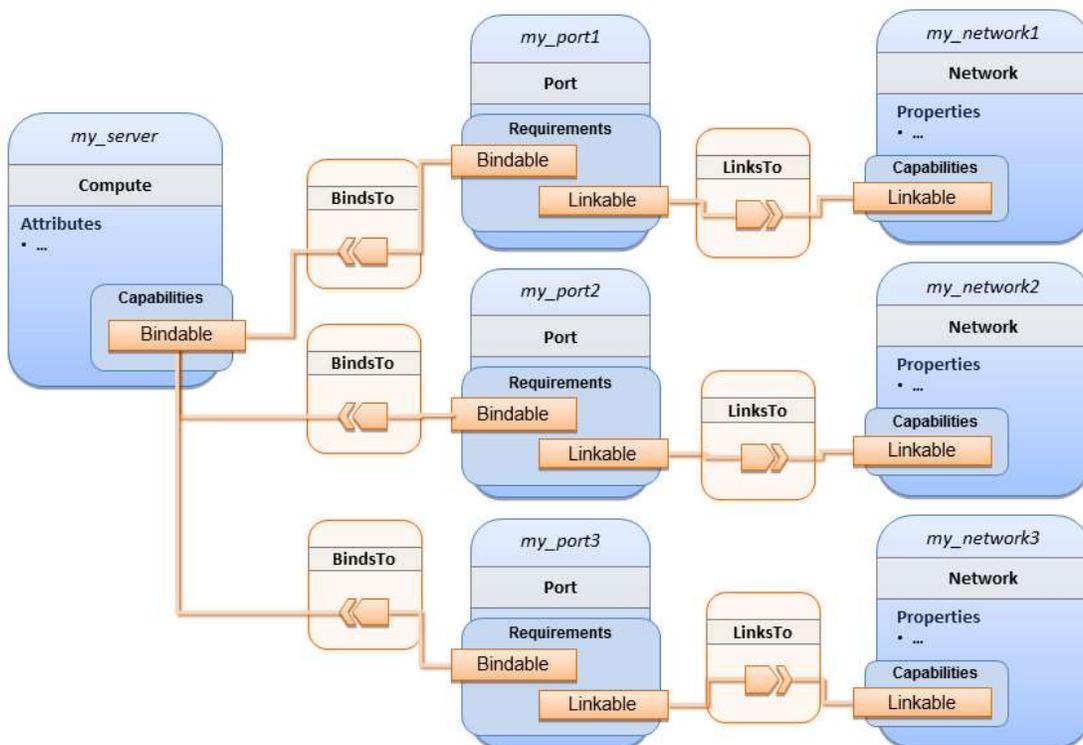
```

2825 **H.1.13 Network 4: Server bound to three networks**

2826 **H.1.13.1 Description**

2827 This use case shows how three logical networks (Network), each with its own IP address range, can be bound to
 2828 with the same server (Compute node).

2829 **H.1.13.2 Logical Diagram**



2830

2831 **H.1.13.3 Sample YAML**

```

tosca_definitions_version: toska_simple_yaml_1_0_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: 512 MB
          os:
            properties:
              architecture: x86_64
              type: Linux
              distribution: CirrOS
              version: 0.3.2

    my_network1:
      type: toska.nodes.network.Network
      properties:
        cidr: '192.168.1.0/24'
        network_name: net1

    my_network2:
      type: toska.nodes.network.Network
      properties:
        cidr: '192.168.2.0/24'
        network_name: net2

    my_network3:
      type: toska.nodes.network.Network
      properties:
        cidr: '192.168.3.0/24'
        network_name: net3

    my_port1:
      type: toska.nodes.network.Port
      properties:
        order: 0
      requirements:
```

```
- binding: my_server
- link: my_network1

my_port2:
  type: toska.nodes.network.Port
  properties:
    order: 1
  requirements:
    - binding: my_server
    - link: my_network2

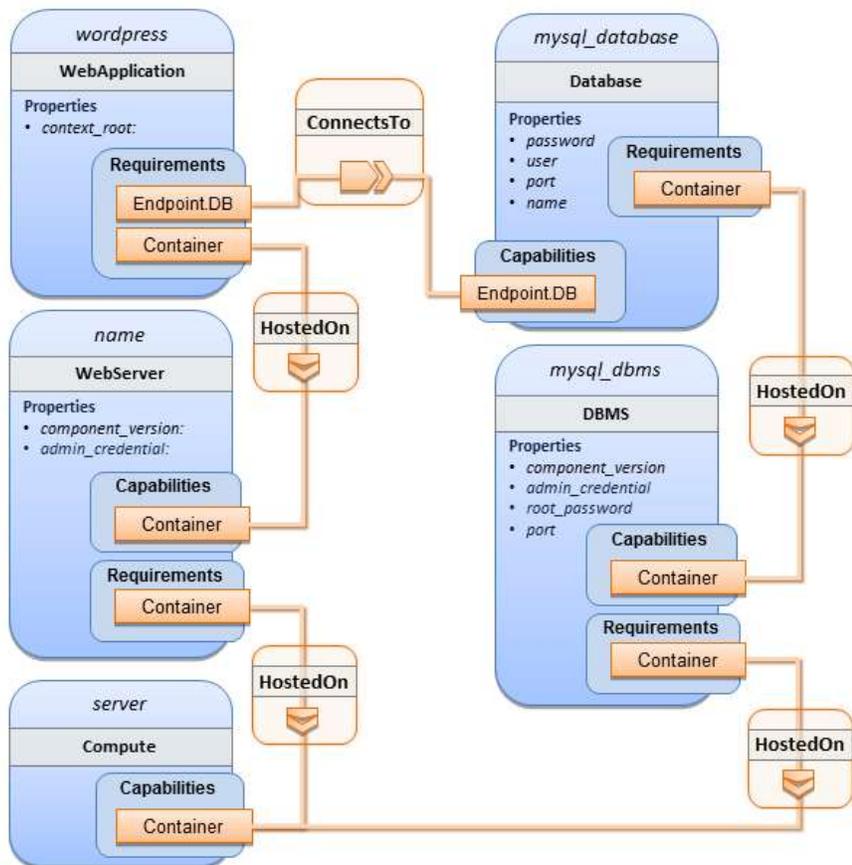
my_port3:
  type: toska.nodes.network.Port
  properties:
    order: 2
  requirements:
    - binding: my_server
    - link: my_network3
```

2832 **H.1.14 WebServer-DBMS 1: WordPress + MySQL, single instance**

2833 **H.1.14.1 Description**

2834 TOSCA simple profile service showing the WordPress web application with a MySQL database hosted on a single
2835 server (instance).

2836 **H.1.14.2 Logical Diagram**



2837

2838 **H.1.14.3 Sample YAML**

```

tosca_definitions_version: tosca_simple_yaml_1_0_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: toasca.nodes.WebApplication.WordPress
    properties:
      context_root: { get_input: context_root }
    requirements:
      - host: webservier
      - 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:
    postconfigure: mysql\_database\_postconfigure.sh

mysql_dbms:
  type: DBMS
  properties:
    root_password: { get_input: db_root_pwd }
    port: { get_input: db_port }
  requirements:
    - host: server
  interfaces:
    Standard:
      create: mysql\_dbms\_install.sh
      start: mysql\_dbms\_start.sh
      configure:
        implementation: mysql\_dbms\_configure.sh
      inputs:
        db_root_password: { get_property: [ mysql_dbms, root_password ] }

webservers:
  type: WebServer
  requirements:
    - host: server
  interfaces:
    Standard:
      create: webservers\_install.sh
      start: webservers\_start.sh

server:
  type: Compute
  capabilities:
    host:
      properties:
        disk_size: 10 GB
        num_cpus: { get_input: cpus }
        mem_size: 4096 kB
    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] }
```

2839 **H.1.14.4 Sample scripts**

2840 Where the referenced implementation scripts in the example above would have the following contents

2841 ***H.1.14.4.1 wordpress_install.sh***

```
yum -y install wordpress
```

2842 ***H.1.14.4.2 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
```

2843 ***H.1.14.4.3 mysql_database_postconfigure.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
```

2844 ***H.1.14.4.4 mysql_dbms_install.sh***

```
yum -y install mysql mysql-server
# Use systemd to start MySQL server at system boot time
systemctl enable mysqld.service
```

2845 ***H.1.14.4.5 mysql_dbms_start.sh***

```
# Start the MySQL service (NOTE: may already be started at image boot time)
```

```
systemctl start mysqld.service
```

2846 **H.1.14.4.6 mysql_dbms_configure**

```
# Set the MySQL server root password  
mysqladmin -u root password db_root_password
```

2847 **H.1.14.4.7 webserver_install.sh**

```
yum -y install httpd  
systemctl enable httpd.service
```

2848 **H.1.14.4.8 webserver_start.sh**

```
# Start the httpd service (NOTE: may already be started at image boot time)  
systemctl start httpd.service
```

2849 **H.1.15 WebServer-DBMS 2: WordPress + MySQL + Floating IPs, single instance**

2850 **H.1.15.1 Description**

2851 This use case is based upon OpenStack Heat's Cloud Formation (CFN) template:

- 2852 • [https://github.com/openstack/heat-](https://github.com/openstack/heat-templates/blob/master/cfn/F17/WordPress_Single_Instance_With_EIP.template)
2853 [templates/blob/master/cfn/F17/WordPress_Single_Instance_With_EIP.template](https://github.com/openstack/heat-templates/blob/master/cfn/F17/WordPress_Single_Instance_With_EIP.template)

2854 **Note:** Future drafts of this specification will detail this use case.

2855 **H.1.15.2 Logical Diagram**

2856 TBD

2857 **H.1.15.3 Sample YAML**

```
TBD
```

2858 **H.1.15.4 Notes**

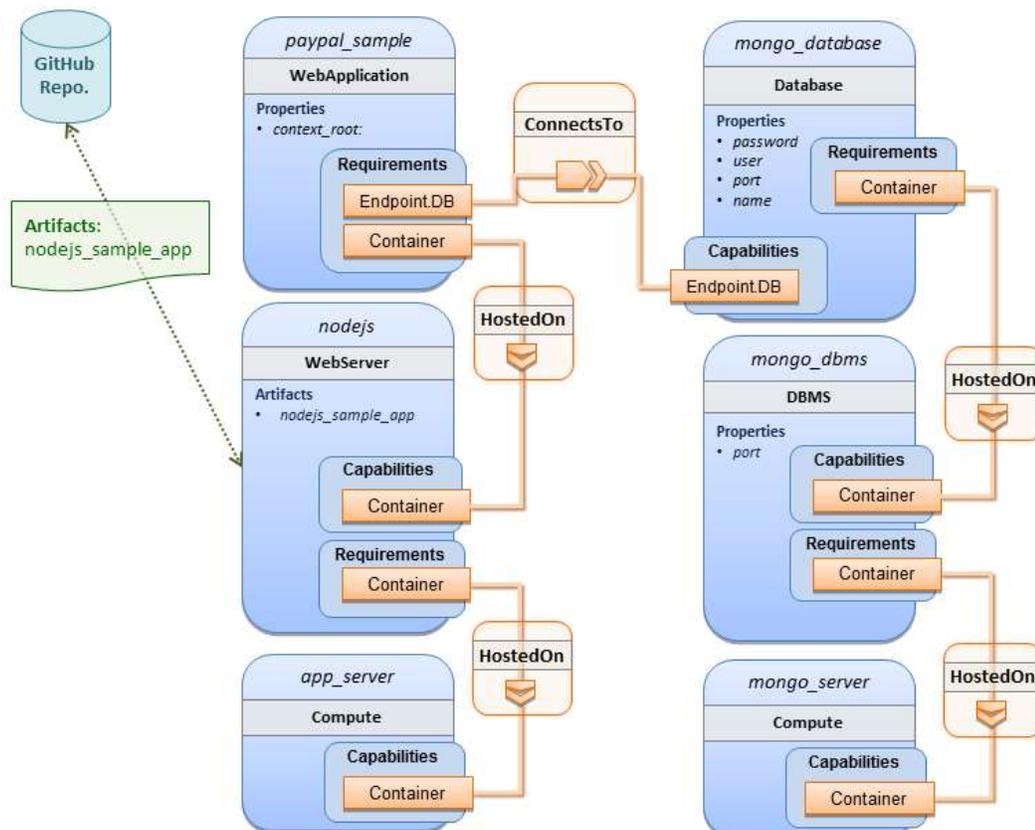
- 2859 • The Heat/CFN use case also introduces the concept of “Elastic IP” (EIP) addresses which is the Amazon
2860 AWS term for floating IPs.
- 2861 • The Heat/CFN use case provides a “key_name” as input which we will not attempt to show in this use
2862 case as this is a future security/credential topic.
- 2863 • The Heat/CFN use case assumes that the “image” uses the “yum” installer to install Apache, MySQL and
2864 Wordpress and installs, starts and configures them all in one script (i.e., under Compute). In TOSCA we
2865 represent each of these software components as their own Nodes each with independent scripts.

2866 **H.1.16 WebServer-DBMS 3: Nodejs with PayPal Sample App and MongoDB on**
2867 **separate instances**

2868 **H.1.16.1 Description**

2869 This use case Instantiates a 2-tier application with Nodejs and its (PayPal sample) WebApplication on one tier
2870 which connects a MongoDB database (which stores its application data) using a ConnectsTo relationship.

2871 **H.1.16.2 Logical Diagram**



2872

2873 **H.1.16.3 Sample YAML**

```
tosca_definitions_version: tosca_simple_yaml_1_0_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: 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
  properties:
    port: 27017
  interfaces:
    tosca.interfaces.node.lifecycle.Standard:
      create: mongodb/create.sh
      configure:
        implementation: mongodb/config.sh
      inputs:
        mongodb_ip: { get_attribute: [mongo_server, private_address] }
      start: mongodb/start.sh

mongo_server:
  type: tosca.nodes.Compute
  capabilities:
    os:
      properties: *os_capabilities
    host:
      properties: *ubuntu_node

app_server:
  type: tosca.nodes.Compute
  capabilities:
    os:
      properties: *os_capabilities
    host:

```

```

properties: *ubuntu_node

outputs:
  nodejs_url:
    description: URL for the nodejs server, http://<IP>:3000
    value: { get_attribute: [app_server, private_address] }
  mongodb_url:
    description: URL for the mongodb server.
    value: { get_attribute: [ mongo_server, private_address ] }

```

2874 **H.1.16.4 Notes:**

- 2875
- Scripts referenced in this example are assumed to be placed by the TOSCA orchestrator in the relative directory declared in TOSCA.meta of the TOSCA CSAR file.
- 2876

2877 **H.1.17 Multi-Tier-1: Elasticsearch, Logstash, Kibana (ELK) use case with multiple instances**

2878

2879 **H.1.17.1 Description**

2880 TOSCA simple profile service showing the Nodejs, MongoDB, Elasticsearch, Logstash, Kibana, rsyslog and collectd installed on a different server (instance).

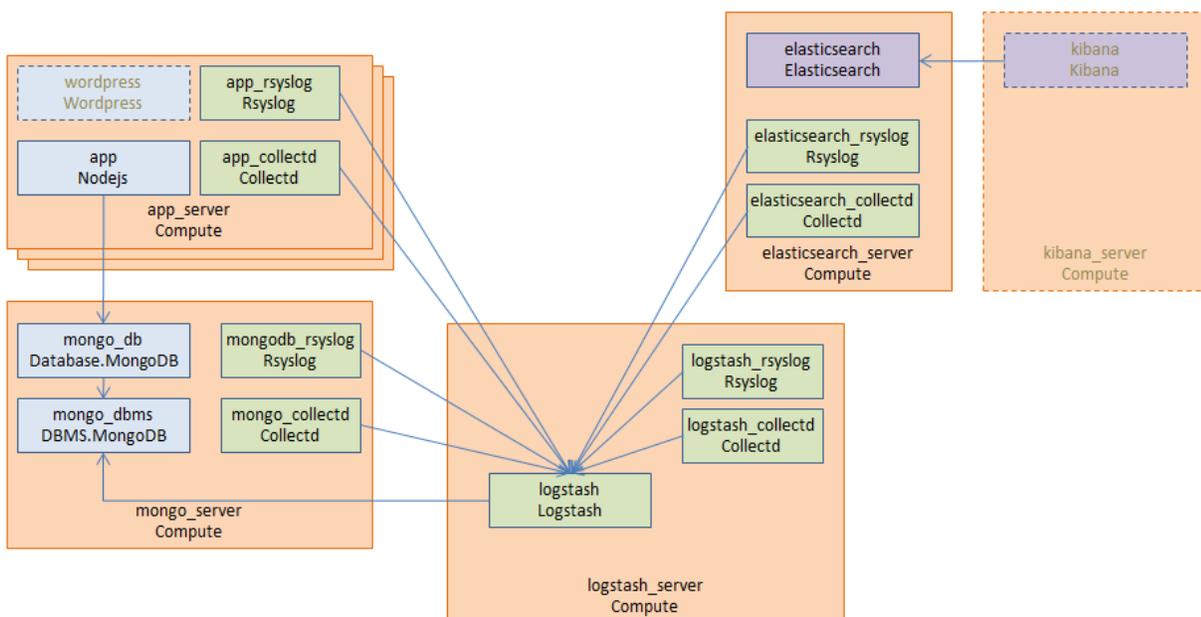
2881

2882

2883 This use case also demonstrates:

- 2884
- Use of TOSCA macros or dsl_definitions
 - Multiple **SoftwareComponents** hosted on same Compute node
 - Multiple tiers communicating to each other over ConnectsTo using Configure interface.
- 2885
- 2886

2887 **H.1.17.2 Logical Diagram**



2888

2889 **H.1.17.3 Master Service Template application (Entry-Definitions)**

2890 TBD

2891

```
tosca_definitions_version: tosca_simple_yaml_1_0_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:
    tosca.interfaces.node.lifecycle.Standard:
      create: Scripts/mongodb/create.sh
      configure:

```

```

    implementation: Scripts/mongodb/config.sh
    inputs:
      mongodb_ip: { get_attribute: [mongo_server, ip_address] }
    start: Scripts/mongodb/start.sh

elasticsearch:
  type: tosca.nodes.SoftwareComponent.Elasticsearch
  requirements:
    - host: elasticsearch_server
  interfaces:
    tosca.interfaces.node.lifecycle.Standard:
      create: Scripts/elasticsearch/create.sh
      start: Scripts/elasticsearch/start.sh
logstash:
  type: tosca.nodes.SoftwareComponent.Logstash
  requirements:
    - host: logstash_server
    - search_endpoint: elasticsearch
  interfaces:
    tosca.interfaces.relationship.Configure:
      pre_configure_source:
        implementation: Python/logstash/configure_elasticsearch.py
        input:
          elasticsearch_ip: { get_attribute: [elasticsearch_server,
ip_address] }
      interfaces:
        tosca.interfaces.node.lifecycle.Standard:
          create: Scripts/logstash/create.sh
          configure: Scripts/logstash/config.sh
          start: Scripts/logstash/start.sh

kibana:
  type: tosca.nodes.SoftwareComponent.Kibana
  requirements:
    - host: kibana_server
    - search_endpoint: elasticsearch
  interfaces:
    tosca.interfaces.node.lifecycle.Standard:
      create: Scripts/kibana/create.sh
      configure:
        implementation: Scripts/kibana/config.sh
        input:
          elasticsearch_ip: { get_attribute: [elasticsearch_server, ip_address]

```

```

}
    kibana_ip: { get_attribute: [kibana_server, ip_address] }
    start: Scripts/kibana/start.sh

app_collectd:
  type: tosa.nodes.SoftwareComponent.Collectd
  requirements:
    - host: app_server
    - collectd_endpoint: logstash
  interfaces:
    tosa.interfaces.relationship.Configure:
      pre_configure_target:
        implementation: Python/logstash/configure_collectd.py
  interfaces:
    tosa.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: tosa.nodes.SoftwareComponent.Rsyslog
  requirements:
    - host: app_server
    - rsyslog_endpoint: logstash
  interfaces:
    tosa.interfaces.relationship.Configure:
      pre_configure_target:
        implementation: Python/logstash/configure_rsyslog.py
  interfaces:
    tosa.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: tosa.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 ] }

```

2892 **H.1.17.4 Sample scripts**

2893 Where the referenced implementation scripts in the example above would have the following contents

2894 **H.1.18 Container-1: Containers using Docker single Compute instance** 2895 **(Containers only)**

2896 *H.1.18.1.1 Description*

2897 This use case shows a minimal description of two Container nodes (only) providing their Docker Requirements
 2898 allowing platform (orchestrator) to select/provide the underlying Docker implementation (Capability). Specifically,
 2899 wordpress and mysql Docker images are referenced from Docker Hub.

2900

2901 This use case also demonstrates:

- 2902 • Abstract description of Requirements (i.e., Container and Docker) allowing platform to dynamically
- 2903 select the appropriate runtime Capabilities that match.
- 2904 • Use of external repository (Docker Hub) to reference image artifact.

2905 **H.1.18.2 Logical Diagram**

2906 TBD

2907 **H.1.18.3 Sample YAML**

2908 *H.1.18.3.1 G1.8.3.1 Two Docker “Container” nodes (Only) with Docker Requirements*

```

tosca_definitions_version: tosca_simple_yaml_1_0_0

description: >
  TOSCA simple profile with wordpress, web server and mysql on the same server.

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.

# Repositories to retrieve code artifacts from
repositories:
  docker_hub: https://registry.hub.docker.com/

topology_template:
  node_templates:

    # The MYSQL container based on official MySQL image in Docker hub
    mysql_container:
      type: tosca.nodes.Container.Application.Docker
      capabilities:
        database_endpoint: tosca.capabilities.Endpoint.Database
      artifacts:
        - my_image: mysql
          type: tosca.artifacts.Deployment.Image.Container.Docker
          repository: docker_hub
      interfaces:
        tosca.interfaces.node.lifecycle.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_endpoint: mysql_container
      artifacts:
        - my_image: wordpress
          type: tosca.artifacts.Deployment.Image.Container.Docker
          repository: docker_hub
      interfaces:
        tosca.interfaces.node.lifecycle.Standard:
          create:
            implementation: my_image
          inputs:

```

```
host_port: { get_input: wp_host_port }
```

2909

2910 Appendix I. Policies (Placeholder)

2911 I.1 Types of policies

2912 Policies typically address two major areas of concern for customer workloads:

- 2913 • Assure governance and compliance with industry and company policies and regulations.
- 2914 • Assure Quality-of-Service and continuity/SLA

2915 I.1.1 SLA policy concerns

- 2916 • Affinity/Anti-affinity: of deployed workloads; that is, what code is allowed to be placed where
- 2917 • Performance (scalability): What resources are an application allowed to consume and

2918 I.1.2 Governance policy concerns

- 2919 • In addition to deploying to a Cloud platform and using a pattern-based approach, customers concerns over
- 2920 “loss of control” are increased. There must be control mechanisms in place that accept governance policies.

2921 I.1.3 Rules considerations

- 2922 • Natural language rules are not realistic, too much to represent in our specification; however, regular
- 2923 expressions can be used that include simple operations and operands that include symbolic names for
- 2924 TOSCA metamodel entities, properties and attributes.
- 2925 • Complex rules can actually be directed to an external policy engine (to check for violation) returns true/false
- 2926 then policy says what to do (trigger or action).
- 2927 • Actions/Triggers could be:
 - 2928 • Autonomic/Platform corrects against user-supplied criteria
 - 2929 • External monitoring service could be utilized to monitor policy rules/conditions against metrics, the
 - 2930 monitoring service could coordinate corrective actions with external services (perhaps Workflow engines
 - 2931 that can analyze the application and interact with the TOSCA instance model).

2932 I.1.4 Definition:

2933 Policies are used to convey a set of capabilities, requirements and general characteristics of an entity.

2934 I.1.5 Policy (Combined requirement)

2935 **Work-in-progress:**

- ```
2936 - name: "my policy"
2937 - type: TBD # categories: affinity (anti-affinity), scaling, performance
```

#### 2937 I.1.6 Requirement (Assertion) Group (grouping construct)

- ```
2938 - "ignorable" | "best can" | "all"
- "choice" (e.g., "one of")
```

2939
2940
2941
2942

I.1.7 Policy Requirement (Assertions)

Use case: Compute1 and Compute2 are 2 node templates. Compute1 has 10 instances, 5 in one region 5 in other region

```
# ----- affinity example -----  
- name: MyAntiAffinityPolicy  
- type: tosa.policy.affinity.  
- rule: fn.separate [ Compute1, Compute2 ] # implies stack-level control  
- trigger: <script>  
  
# ---scaling example ----  
  
- name: MyScaleUpPolicy  
- type: tosa.policy.scale.up | tosa.policy.scale.down  
- rule: fn.utilizaton [ Compute1, Compute2 ], greater_than: 80%  
- trigger: <script>
```

2943

Appendix J. References

2944

J.1 Known Extensions to TOSCA v1.0

2945
2946

The following items will need to be reflected in the TOSCA (XML) specification to allow for isomorphic mapping between the XML and YAML service templates.

2947

J.1.1 Model Changes

2948
2949

- The “TOSCA Simple ‘Hello World’” example introduces this concept in Section 3. Specifically, a VM image assumed to be accessible by the cloud provider.

2950

- Introduce template Input and Output parameters

2951

- The “Template with input and output parameter” example introduces concept in Section 3.1.

2952

- “Inputs” could be mapped to BoundaryDefinitions in TOSCA v1.0. Maybe needs some usability enhancement and better description.

2953

- “outputs” are a new feature.

2954

- “outputs” are a new feature.

2955

- Grouping of Node Templates

2956

- This was part of original TOSCA proposal, but removed early on from v1.0. This allows grouping of node templates that have some type of logically managed together as a group (perhaps to apply a scaling or placement policy).

2957

- This was part of original TOSCA proposal, but removed early on from v1.0. This allows grouping of node templates that have some type of logically managed together as a group (perhaps to apply a scaling or placement policy).

2958

- This was part of original TOSCA proposal, but removed early on from v1.0. This allows grouping of node templates that have some type of logically managed together as a group (perhaps to apply a scaling or placement policy).

2959

- Lifecycle Operation definition independent/separate from Node Types or Relationship types (allows reuse).

2960

For now we added definitions for “node.lifecycle” and “relationship.lifecycle”.

2961

- Override of Interfaces (operations) in the Node Template.

2962

- Service Template Naming/Versioning

2963

- Should include TOSCA spec. (or profile) version number (as part of namespace)

2964

- Allow the referencing artifacts using a URL (e.g., as a property value).

2965

- Repository definitions in Service Template.

2966

- Substitution mappings for Topology template.

2967

J.1.2 Normative Types

2968

- Constraints

2969

- constraint clauses, regex

2970

- Types / Property / Parameters

2971

- list, map, range, scalar-unit types

2972

- Includes YAML intrinsic types

2973

- NetworkInfo, PortInfo, PortDef, PortSpec, Credential

2974

- TOSCA Version based on Maven

2975

- Node

2976

- Root, Compute, ObjectStorage, BlockStorage, Network, Port, SoftwareComponent, WebServer, WebApplicaton, DBMS, Database, Container, and others

2977

- Root, Compute, ObjectStorage, BlockStorage, Network, Port, SoftwareComponent, WebServer, WebApplicaton, DBMS, Database, Container, and others

2978

- Relationship

2979

- Root, DependsOn, HostedOn, ConnectsTo, AttachesTo, RoutesTo, BindsTo, LinksTo and others

2980

- Artifact

2981

- Deployment: Image Types (e.g., VM, Container), ZIP, TAR, etc.

2982

- Implementation: File, Bash, Python, etc.

2983

- Requirements

2984

- None

2985

- Capabilities

- 2986 • Container, Endpoint, Attachment, Scalable, ...
- 2987 • Lifecycle
- 2988 • Standard (for Node Types)
- 2989 • Configure (for Relationship Types)
- 2990 • Functions
- 2991 • get_input, get_attribute, get_property, get_nodes_of_type, get_operation_output and others
- 2992 • concat, token
- 2993 • get_artifact

2994 J.2 Terminology

2995 The key words “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, “SHOULD”, “SHOULD NOT”,
 2996 “RECOMMENDED”, “MAY”, and “OPTIONAL” in this document are to be interpreted as described in

[TOSCA-1.0]	Topology and Orchestration Topology and Orchestration Specification for Cloud Applications (TOSCA) Version 1.0, an OASIS Standard, 25 November 2013, http://docs.oasis-open.org/tosca/TOSCA/v1.0/os/TOSCA-v1.0-os.pdf
[YAML-1.2]	YAML, Version 1.2, 3rd Edition, Patched at 2009-10-01, Oren Ben-Kiki, Clark Evans, Ingy döt Net http://www.yaml.org/spec/1.2/spec.html
[YAML-TS-1.1]	Timestamp Language-Independent Type for YAML Version 1.1, Working Draft 2005-01-18, http://yaml.org/type/timestamp.html

2997 .

2998 J.3 Normative References

[TOSCA-1.0]	Topology and Orchestration Topology and Orchestration Specification for Cloud Applications (TOSCA) Version 1.0, an OASIS Standard, 25 November 2013, http://docs.oasis-open.org/tosca/TOSCA/v1.0/os/TOSCA-v1.0-os.pdf
[YAML-1.2]	YAML, Version 1.2, 3rd Edition, Patched at 2009-10-01, Oren Ben-Kiki, Clark Evans, Ingy döt Net http://www.yaml.org/spec/1.2/spec.html
[YAML-TS-1.1]	Timestamp Language-Independent Type for YAML Version 1.1, Working Draft 2005-01-18, http://yaml.org/type/timestamp.html

2999 J.4 Non-Normative References

[AWS-CFN]	Amazon Cloud Formation (CFN), http://aws.amazon.com/cloudformation/
[Chef]	Chef, https://wiki.opscode.com/display/chef/Home
[OS-Heat]	OpenStack Project Heat, https://wiki.openstack.org/wiki/Heat
[Puppet]	Puppet, http://puppetlabs.com/
[WordPress]	WordPress, https://wordpress.org/
[Maven-Version]	Apache Maven version policy draft: https://cwiki.apache.org/confluence/display/MAVEN/Version+number+policy

3000 J.5 Glossary

3001 The following terms are used throughout this specification and have the following definitions when used in
 3002 context of this document.

Term	Definition
Instance Model	A deployed service is a running instance of a Service Template. More precisely, the instance is derived by instantiating the Topology Template of its Service

Template, most often by running a special plan defined for the Service Template, often referred to as build plan.

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

Appendix K. Issues List

Issue #	Target	Status	Owner	Title	Notes
TOSCA-132	CSD05	Defer	Palma	Use "set_property" methods to "push" values from template inputs to nodes	Feature. Needs new owner.
TOSCA-135	CSD04	Open	Palma	Define/reference a Regex language (or subset) we wish to support for constraints	Feature, Reference a Perl subset.
TOSCA-136	CSD04	Open	Zala	Need rules to assure non-collision (uniqueness) of requirement or capability names	None
TOSCA-137	CSD05	Defer	Palma	Need to address "optional" and "best can" on node requirements (constraints) for matching/resolution	
TOSCA-138	CSD05	Defer	Palma	Define a Network topology for L2 Networks along with support for Gateways, Subnets, Floating IPs and Routers	Luc Boutier has rough proposal in MS Word format.
TOSCA-140	CSD04	Open	Palma	Constraining the capabilities of multiple node templates	
TOSCA-141	CSD04	Open	Palma	Specifying Environment Constraints for Node Templates (Policy related)	
TOSCA-142	CSD03	Open Fixed	Spatzier / Rutkowski	Define normative Artifact Types (including deployment/packages, impls., and runtime types)	
TOSCA-143	CSD03	Open Fixed	Rutkowski	Define normative tosca.nodes.Network Node Type (for simple networks)	Separate use case as what Luc proposes in TOSCA-138.
TOSCA-148	CSD03	Open Fixed	Palma	Need a means to express cardinality on relationships (e.g., number of connections allowed)	Occurrences now on capability and requirement defs.
TOSCA-151	CSD04	Defer	Rutkowski	Resolve spec. behavior if name collisions occur on named Requirements	subtask of TOSCA-148
TOSCA-152	CSD05	Defer	Palma	Extend Requirement grammar to support "Optional/Best Can" Capability Type matching	subtask of TOSCA-137
TOSCA-153	CSD04	Open	Rutkowski	Define grammar and usage of Service Template keyname (schema namespace) "tosca_default_namespace"	
TOSCA-154	CSD05	Defer	Palma	Decide how security/access control work with Nodes, update grammar, author descriptive text/examples	Deferred
TOSCA-155	CSD05	Defer	Rutkowski	How do we provide constraints on properties declared as simple YAML lists (sets)	Deferred
TOSCA-156	CSD05	Defer	Palma	Are there IPv6 considerations (e.g., new properties) for tosca.capabilities.Endpoint	Deferred
TOSCA-158	CSD05	Defer		Provide prose describing how Feature matching is done by orchestrators	Deferred. Subtask of TOSCA-137
TOSCA-161	CSD03	FIXED	Spatzier	Need examples of using the built-in feature (Capability) and dependency (Requirement) of tosca.nodes.Root	Deferred; however the Root node was fixed and the Root capability type was added
TOSCA-162	CSD05	Defer	Rutkowski	Provide recognized values for tosca.nodes.compute properties: os_arch	Deferred
TOSCA-163	CSD05	Defer	Vachnis	Provide recognized values for tosca.nodes.BlockStorage: store_fs_type	Deferred
TOSCA-165	CSD05	Defer	Need new owner	New use case / example: Selection/Replacement of web server type (e.g. Apache, NGinx, Lighttpd, etc.)	Deferred
TOSCA-166	CSD05	Defer	Unassigned	New use case / example: Web Server with (one or more) runtimes environments (e.g., PHP, Java, etc.)	Deferred

TOSCA-167	CSD05	Defer	Unassigned	New use case / example: Show abstract substitution of Compute node OS with different Node Type Impls.	Deferred
TOSCA-168	CSD06	Defer	Unassigned	New use case / example: Show how substitution of IaaS can be accomplished.	Deferred
TOSCA-170	CSD05	Defer	Elisha	WD02 - Explicit textual mention, and grammar support, for adding (extending) node operations	Deferred
TOSCA-172	CSD05	Defer	Lipton	2014 March - Public Comment Questions (Plans, Instance Counts, and linking SW Nodes)	Deferred
TOSCA-176	CSD03	Fixed	Elisha	Add connectivity ability to Compute	Deferred. However, we added Endpoint.Admin to Compute. This is intended for SSH (using ConnectsTo)
TOSCA-179	CSD03	Defer CLOSE	Elisha	Add "timeout" and "retry" keynames to an operation	Deferred. However, we do NOT intend to have TOSCA define how "retry" on connections should be implemented. Recommend Close.
TOSCA-180	CSD02	Open / In- progress FIXED	Rutkowski	Support of secured repositories for artifacts	Repository support added.
TOSCA-181	CSD03	Open FIXED	Boutier	Dependency requirement type should match any target node.	Subtask of TOSCA-161 161 is fixed and now any node derived from Root node can require any other node also derived from Root
TOSCA-182	CSD04	Defer	Palma	Document parsing conventions	Deferred
TOSCA-183	CSD04	Open	Palma	Composition across multiple yaml documents	Deferred
TOSCA-184	CSD05	Defer	Palma	Pushing (vs pulling) inputs to templates	Subtask of TOSCA-132, Deferred
TOSCA-185	CSD05	Defer	Durand	Instance model	Deferred
TOSCA-186	CSD04	Defer	Spatzier	model composition	Deferred
TOSCA-189	CSD03	Open	Shtilman	Application Monitoring - Proposal	Monitoring WG should use as a use case / discussion
TOSCA-191	CSD03	Open Fixed	Rutkowski	Document the "augmentation" behavior after relationship is selected in a requirement	
TOSCA-193	CSD04	Open Review	Spatzier	"implements" keyword needs its own section/grammar/example in A.5.2	Subtask of TOSCA-186 See if we can close this as much as been fixed, but perhaps a small feature remains and deserves its own issue
TOSCA-194	CSD02	Open Review	Lauwers	Nested Service Templates should be able to define additional operations	Subtask of TOSCA-186
TOSCA-200	CSD04	Open Review	Vachnis, Parasol	Query based upon capability	New instance model functions to be provided.

TOSCA-201	CSD03	Deferred	Lauwers	Harmonize Properties and Capabilities in Node Types	Deferred
TOSCA-202	CSD03	Open FIXED	Boutier	Cardinalities for capabilities and requirements	Subtask of TOSCA-148 Verify fix with Luc.
TOSCA-205	CSD03	Open	Boutier	Add interface type.	
TOSCA-208	CSD03	Open	Boutier	Add conditional capabilities (enable/disable capabilities on a node)	
TOSCA-209	CSD02	Open	Rutkowski	Fix Grouping example to use correct parameter for WebServer	
TOSCA-210	CSD02	Open	Rutkowski	Need example on get_xxx functions using HOST keyword	
TOSCA-211	CSD02	Open	Rutkowski	Need version on TOSCA Types (Node, Relationship, etc.)	
TOSCA-213	CSD03	Open	Lauwers	Clarify distinction between declaring properties and assigning property values	
TOSCA-214	CSD02	Open	Vachnis / Rutkowski	New functions for accessing the instance model	
TOSCA-217	CSD02	Open	Spatzier / Rutkowski	Add new simplified, single-line list notation / grammar for Requirement Def.	
TOSCA-219	CSD03	Open	Boutier	Workflow/Plan generation and components state dependency	
TOSCA-220	CSD03	Open	Boutier	get_artifact function	Proposal in DOCX lined to issue
TOSCA-224	CSD03	Open	Rutkowski	Discuss removing "github_url" property from non-normative Nodejs Node Type	
TOSCA-225	CSD03	New	Palma	Adding a (untyped) dependency relationship between Node Templates	
TOSCA-226	CSD03	New	Spatzier	Proposal (incomplete) to change Requirement Def. from Ordered List to a Map	
TOSCA-227	CSD03	New	Spatzier	Proposal (incomplete) to change Property Filters Ordered List to a Map	
TOSCA-228	CSD05	Defer	Rutkowski	Discuss "Final" concept for Types, Operation and properties	Matt agrees to Defer
TOSCA-229	CSD03	Open	Boutier	Requirements / target filters and subsituable / selectable are too confusing.	
TOSCA-230	CSD03	Open	Zala	TOSCA artifacts doesn't resolve version dependency	
TOSCA-231	CSD03	Open	Boutier	Scalar units and get_property	
TOSCA-232	CSD03 Fixed	Open	Shitao Li	Adding metadata section in Service template	NFV WG submitted
TOSCA-233	CSD03	Open	Shitao Li	How to handle VM image in terms of software component	We now know how we want to do this, just need example in Appendix H to show
TOSCA-234	CSD03	Open	Rutkowski	Discuss removing "github_url" property from non-normative Nodejs Node Type	
TOSCA-235	CSD03	Open	Rutkowski	Endpoint capability needs support for Fixed IPs	Relates to Floating IP work.
TOSCA-236	CSD03	Open Close	Rutkowski	Mark IaaS Node Types as "substitutable" since the Orchestrator can determine impl. by default	No longer valid, all nodes are now "selectable" can be substituted with another topology (template).
TOSCA-238	CSD03	Open	Zala, Rutkowski	Task: Author Section A.1.1 "A.1.1 Rules to avoid namespace collisions	
TOSCA-239	CSD03	Open CLOSED	Rutkowski	Feature: Address override comment in Chapter 5	
TOSCA-240	CSD03	Open	Rutkowski	Ch6. How are artifacts referenced from CSAR file?	

				Automatic copy/delete?	
TOSCA-241	CSD04	Open	Rutkowski	Task: See if OASIS will allow us to use "fix" version in TOSCA version strings	
TOSCA-242	CSD04	Open	[Zala]	How to use 'output' section - new use case needed	
TOSCA-243	CSD04	Open	[Zala]	CSD04: Need mechanism to retrieve parent type's operation script	
TOSCA-244	CSD04	Open	Boutier	Need to add clarity around orchestrator's copy behavior for script	
TOSCA-245	CSD04	Open	[Zala]	Need clear description of relationship priority/ordering during deployment	
TOSCA-246	CSD04	Open	[Zala]	Define global/environment dependencies emphasis on script processors	
TOSCA-247	CSD03	Open Fixed	Rutkowski	Need description of abstract/subst. node types and application of node_filters	
TOSCA-248	CSD04	Open	Spatzier	Consider mapping inputs and outputs and show different naming with example	
TOSCA-249	CSD04	Open	Rutkowski	Verify artifacts definition grammar applies to both artifacts type and artifacts templates	
TOSCA-250	CSD04	Open	Rutkowski	Determine how to avoid collision of artifacts types installed into node environments	
TOSCA-251	CSD04	Open	[Zala]	Explore moving Appendix G (modeling use cases) to earlier chapters.	

3005

3006

Appendix L. Acknowledgments

3007 The following individuals have participated in the creation of this specification and are gratefully acknowledged:

3008 **Contributors:**

- 3009 Avi Vachnis (avi.vachnis@alcatel-lucent.com), Alcatel-Lucent
3010 Chris Lauwers (lauwers@ubicity.com)
3011 Derek Palma (dpalma@vnomnic.com), Vnomic
3012 Frank Leymann (Frank.Leymann@informatik.uni-stuttgart.de), Univ. of Stuttgart
3013 Gerd Breiter (gbreiter@de.ibm.com), IBM
3014 Hemal Surti (hsurti@cisco.com), Cisco
3015 Ifat Afek (ifat.afek@alcatel-lucent.com), Alcatel-Lucent
3016 Idan Moyal, (idan@gigaspace.com), Gigaspaces
3017 Jacques Durand (jdurand@us.fujitsu.com), Fujitsu
3018 Jin Qin, (chin.qinjin@huawei.com), Huawei
3019 Juergen Meynert (juergen.meynert@ts.fujitsu.com), Fujitsu
3020 Kapil Thangavelu (kapil.thangavelu@canonical.com), Canonical
3021 Karsten Beins (karsten.beins@ts.fujitsu.com), Fujitsu
3022 Kevin Wilson (kevin.l.wilson@hp.com), HP
3023 Krishna Raman (kraman@redhat.com), Red Hat
3024 Luc Boutier (luc.boutier@fastconnect.fr), FastConnect
3025 Matt Rutkowski (mrutkows@us.ibm.com), IBM
3026 Moshe Elisha (moshe.elisha@alcatel-lucent.com), Alcatel-Lucent
3027 Nate Finch (nate.finch@canonical.com), Canonical
3028 Nikunj Nemani (nnemani@vmware.com), VMware
3029 Richard Probst (richard.probst@sap.com), SAP AG
3030 Sahdev Zala (spzala@us.ibm.com), IBM
3031 Shitao li (lishitao@huawei.com), Huawei
3032 Simeon Monov (sdmonov@us.ibm.com), IBM
3033 Stephane Maes (stephane.maes@hp.com), HP
3034 Thomas Spatzier (thomas.spatzier@de.ibm.com), IBM
3035 Ton Ngo (ton@us.ibm.com), IBM
3036 Travis Tripp (travis.tripp@hp.com), HP
3037 **Vahid Hashemian** (vahidhashemian@us.ibm.com), IBM
3038 Wayne Witzel (wayne.witzel@canonical.com), Canonical
3039 Yaron Parasol (aronpa@gigaspace.com), Gigaspaces

Appendix M. Revision History

Revision	Date	Editor	Changes Made
WD05, Rev01	2014-12-11	Matt Rutkowski, IBM	<ul style="list-style-type: none"> Initial WD05, Revision 01 baseline.
WD05, Rev02	2015-01-18	Matt Rutkowski, IBM	<ul style="list-style-type: none"> Separated out scalar-unit.size while adding base 2 values and added scalar-unit.time with examples and informative references. Fixed constraints in Endpoint to allow for either (optionally) port as PortDef or map of PortSpec. Updated properties using old scalar-unit to use scalar-unit.size
WD05, Rev03	2015-01-21	Matt Rutkowski, IBM	<ul style="list-style-type: none"> Fixed missing requires=false clauses on port and ports properties of Endpoint. Removed anonymous types for lists and maps. Removed reference from schema definition form Properties and adjusted Properties to reflect resulting set of reduced properties for entry0-schema. Incorporated comments from Chris L. addressed/fixed typos and non-material corrections identified. Fixed PortSpec definition and example from Victor. Fixed other typos found while integrating Chris' comments.
WD05, Rev04	2015-01-29	Matt Rutkowski, IBM	<ul style="list-style-type: none"> Removed Simple interface to prevent invalid YAML on the TOSCA Root node type and avoids describing multiple instance/sequencing of methods with different names and would cause type proliferation (standard vs. simple types). Section 18 example: Fixed invalid YAML and assured TOSCA-191 has correct problem stated. A.4.14, C.2.3 Removed the "type" keyword from schema (datatype) definitions as there should be only one keyword (i.e., "derived_from") used to define new types. It also made no sense to continue existing paradigm of using "derived_from" as we have in all other places. A.4.2.1: Added range type to in_range constraint and map and list types to some length-based constraints. Added scalar-unit types to list of comparable types with requirements to include units in comparison. A.4.4: The "type" keyname on a Property Definition is no longer mandatory as new datatypes are not required to derive from an existing type. Formerly A.4.14: Removed "type" keyname from schema in favor of "derived_from" as both were not needed since data_types were added. A.4.14, A.4.15: merged "Schema Definition" from A4.14 into "Datatype definition" as this is the only place this grammar is referenced since removing it from "entry_schema" grammar. A.4.21: Fixed unclear use/description of property filters under named capabilities listed on the overall filter definition. A.4.21: Added "attributes" keyname (and attribute definitions) to the Capability Type. A.4.26: Added Attribute value assignment section (A4.26) A.4.27: Added Property value assignment section (A4.26) A.5.1: Changed "datatype_definitions" to just "data_types" to match the names of other keys used in the service template (e.g., node_types, relationship_types, etc.).
WD05, Rev05	2015-02-02	Matt Rutkowski, IBM	<ul style="list-style-type: none"> A.4.5: documented "status_value" in grammar A.5.2.2 – Fixed map example; missing ContactInfo and copy-paste error. TOSCA-218 - A.4.22, C.3.1.1 – Changed "valid_node_types" from a property

			<p>to a meta-level keyname of “valid_node_types”. This affects the grammar of several normative TOSCA capability types.</p> <ul style="list-style-type: none"> • Merged in numerous additional comments from Chris against WD05, Rev03, mostly in Sections 1-14 but also a few grammar comments in the topology template section. • Port Type requirements “binding” and “link” were not ordered list (fixed) and link should appear before binding. Fixed network example to match. • Added a section A.4.1 to clearly state that required keynames do not have to appear in types that derive from another (parent) type that already provides them. • Changed meta-level keynames “valid_node_types” to the key “valid_types” which allows more than nodes as sources for a capability type. • Added a “Required” column for most of the Keyname tables for Appendix A in order to better see which keynames are required. • A.4.3.2 - Clarified that the “equals” constraint (operand) is optional and if a value simply appears it should be interpreted as “equals”. • Fixed examples in sections 6 and 7 to reference the latest non-normative and normative type properties and methods (e.g., get_attribute for ip_address). • A.4.18, F.1.3 – Identified issues in Requirement definition (A.4.18) as shown use case in F.1.3. Documented as “mustfix”.
WD05, Rev06	2015-02-05	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • Appendix H: Added a placeholder section with early considerations for policies, their grammar and use cases. • Resolved to preserve the description of direction (i.e., source and target) in the name of “valid_types” keynames that indicate what nodes types are allowed on each end of a TOSCA Relationship. <ul style="list-style-type: none"> ○ A.4.25 - In the case of Relationship Types, we will use “valid_target_types” to indicate what types a relationship can be connected to (point to or target). ○ A.4.23 - For Capability Types we use “valid_source_types” to indicate what source node types are allowed to form relationships to the nodes the capability are declared in. ○ Adjusted all examples to use new keynames. • C.5.1 – The key “valid_target_type” for toska.relationships.Root should not have been included and was removed. It had indicated that any toska.nodes.Root. • F.1.4 – Removed duplication of normative type definitions included in modeling use case /example. This caused us to maintain normative types accurate in two places and was confusing to readers. • G.1.4.4, G.1.4.5 – Moved custom AttachesTo relationship from x.5 to x.4 and adjusted to use current grammar. • E.5.2 - Fixed toska.nodes.network.Port definition to use correct “node” keyword • E.6 - Adjusted example templates in E.6.1 and E.6.2 to change “tosca.nodes.Port” to “tosca.nodes.network.Port” for all occurrences. • E.6 - Adjusted example templates in E.6.1 and E.6.2 to change “tosca.nodes.Network” to “tosca.nodes.network.Network” for all occurrences. • A.2.3: Added “UNBOUNDED” keyword to TOSCA range type • A.4.18: Added “occurrences” keyname to Requirement def. • C.3.2, C.5.2: Added “tosca.capabilities.node” and added it as the default target type for the DependsOn relationship.

			<ul style="list-style-type: none"> • C.7.1: Fixed “tosca.nodes.root definition” <ul style="list-style-type: none"> ○ Added “state” attribute def. to grammar ○ Fixed the built-in “dependency” requirement adding cardinality default of [0, UNBOUNDED]. This allows section 10 example to remain valid AND fixes all normative node type definitions. ○ Added “feature” capability of type “tosca.capabilities.node” • C.5.2.1: Added the tosca.capabilities.Root to “valid_target_types” of the DependsOn Relationship type. • Section 13.2, Section 14: Merged in template “substitution” use cases in place of old “nested templates” to reflect current WG focus in this area. • C.5.3.1: HostedOn now derives from tosca.relationships.Root. • A.3.2: Directives section to describe reserved values for “directives” keyname and added “substitutable” value for Node Templates. • C.5.3: Changed HostedOn to derive from Root instead of DependsOn • D.2.4: Changed “database_endpoint” requirement to “wordpress_database_requirement” and fixed it have the ConnectsTo relationship and “database_endpoint” capability values. This will allow us to test the ability to connect a relationship to a type-compatible capability in another node when the symbolic names do not match. Updated use cases for wordpress. • Appendix A: Reorganized section to group reusable element, types and template definitions together. Also removed simple “list” element definitions as they served little purpose and moved their grammars in-line to where they were referenced. • A.6.4 - Added the Interface Type to allow new interfaces to be defined; it is slightly different than an interface definition (more restrictive). Please read “Additional Requirements”. Also, added a means to define them in the Service Template using the interfaces keyword. • C.5.1, C.7.1 - Changed the “interfaces” keyword grammar for both tosca.nodes.Root and tosca.relationships.Root to reflect the correct way to declare an interface (not using a simple list). • A.5.9 - Added “type” keyword to Interfaces definition and changed example to better distinguish it from Interface Type example. • A.6 Service Template - Added note describing how Types can be defined in Service Templates (no topology-template) for import/use in other service templates. • C.3.5 - Changed tosca.capabilities.DatabaseEndpoint to tosca.capabilities.EndPoint.Database. Updated MySQL Node Type to reflect this change. We do not need to specify an Endpoint.Database.MySQL capability type as the requirement should differentiate the database by its Node Type (MySQL) and not its capability. • C.7.2: added “local_storage” requirement (occurrences [0,unbounded]) to actually allow us a built-in way to attach BlockStorage nodes. • C.3.5: Added tosca.capabilities.Endpoint.Admin • C.5.4: Added Credential property to ConnectsTo relationship
WD05, Rev07	2015-03-02	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • A.5.7: Added the “final” keyname to the Property definition. Added final: true to the Endpoint.Admin capability property “secure” so subclasses cannot change this value. • Clarified C.5.5 AttachesTo relationship properties. • B.4.2, B.5.1 – Fixed get_attribute and get_property grammar to not confuse the square bracket '[' meaning “optional parameter” with the YAML square bracket for simple list (set).

			<ul style="list-style-type: none"> • B.3, B.3.1: Added “Intrinsic function section and the definition of the “concat” function to address TOSCA-212. • F.1.1: Fixed “valid_source_types” still listed as a property and corrected description of its use in the prose • A.6.1: Assured valid_source_types” was in Capability definition. • C.7.10: Containee Node Type placeholder.
WD05, Rev08	2015-03-03	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • A.5.7: Removed “final” from properties until we have a comprehensive discussion on a “final” designator for Type definitions (Nodes, Rels. Data, etc.) and on interfaces/operations (to lock down implementation scripts). <ul style="list-style-type: none"> ◦ Opened JIRA issue TOSCA-228 to discuss • A.5.8: Fixed Operation definition to include different grammars for use in Types versus Templates (Node and Relationship). Also provided new Additional requirements that describe artifact override behaviors. • A.5.8: Fixed Operation definition to include different grammars for use in Types versus Templates (Node and Relationship). • A.6.2: Removed “node_filter” (keyword and grammar) from Node Type requirement definitions as it should only be valid on Node Template definitions. • A.7.4: Node Template: Added the “node_filter” keyname (with node filter grammar) and additional requirements on usage in conjunction with the “selectable” directive. • A.3.2: Added “selectable” directive keyword. • A.6.1: Added valid_source_types to Capability definition and added a requirement that only allows for subclasses to provide Node Types (names) for this keyname that are type-compatible and derived from types declared in their parent class. • 11.1: Authored a section to show how “selectable” can be used as an alternative way to provide dynamic matching to what is described earlier in Section 11. • Assure the consistent use of “keyname” versus “keyword” throughout the document. • 17.1: Removed Use case: “Providing input properties for all interfaces” which has no comparison in other language designs. In addition, we had no real or practical examples for this nor had we added the grammar for this. • 19: Fixed example to use a relationship template, previously we were using get_attribute to retrieve a value on the abstract ConnectsTo relationship directly on a Node Type definition which is illegal in the grammar. • 8: Fixed example to use a relationship template for a custom connection as the grammar does not allow behaviors (scripts) to be provided on abstract ConnectsTo relationship types on Requirement definitions. • F.1.4: Removed invalid use of “default” keyname on Relationship Type examples. • F.1.5: Removed placeholder (empty) section for examples of “add_target”, “remove_target” and “target_changed” operations as we should explore these as part of then-tier logging use case in Section G.1.9.
WD05, Rev09	2015-03-17	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • Sections 1-19: added “highlighting” to keywords in code samples that were being featured in the surrounding text (to better highlight or feature the new idea being introduced by that section). • A.7.4, A.7.5: Added the “copy” keyname to Node and Relationship Template definitions / grammars. This was described in use case in section F.1.4 which was updated to reflect the name change to “copy” from “alias” which was not a good descriptive keyname.

			<ul style="list-style-type: none"> • B.2.1: Clarified where the IDs come from (i.e., <code>tosca_id</code> attribute). • C.3.8.1: <code>OperatingSystem</code> Capability type, changed type and architecture to optional from required allowing more nodes to use this capability. • A2, C.7: Added the type name (i.e., <code>shortname</code>, <code>qualified</code> and <code>typeURI</code>) for every TOSCA datatype that did not have one declare. • A.2.2: Added constraints for TOSCA version type • C.7.6, C.7.7: Removed prefixes on properties to follow working group decision to not use prefixes and have already removed them on other node types. This means prefixes <code>"db_"</code> and <code>"dbms_"</code> are removed from property names on Database and DBMS Node Types respectively. • C.7.8: Remove <code>"store_"</code> prefix from <code>ObjectStore</code> • C.7.9: Removed redundant attribute <code>"Volumelid"</code> as we have <code>"volume_id"</code> as a property which would be the (reflected) attribute we would use to retrieve the ID assigned by the platform. • D.1.4: removed <code>"db_"</code> prefix from non-normative <code>WordPress</code> node type. • A.8.1: Added example for relationship template • G.1.9: Renamed the <code>"N-Tier"</code> use case to <code>Elasticsearch</code>, <code>Logstash</code>, <code>Kibana</code> (ELK) use case to better feature the new open source components the use case is adding. • A.5.8: Removed <code>"default"</code> keyname from attribute. • A.5.9: Added clarification (additional requirement) that allows implementations to provide attribute values at runtime apart from template-included value expressions. • C.7.2, C.7.3: Removed deprecated <code>ip_address</code> property from <code>Compute</code> and <code>SoftwareComponent</code> node types. This was used for demos in early drafts and should not be in v1.0, especially now that we have a proper network model supporting multiple <code>IPAddresses</code>. We now have clearly distinguishable <code>"public_address"</code> and <code>"private_address"</code> attributes in <code>Compute</code> node. Cascade the name change across all use cases in document. • E.5.2: Removed redundant <code>"ip_address"</code> attribute (since it already is a property it is automatically also an attribute). • C.3.8: Fixed copy-paste error for <code>OperatingSystem</code> <code>"Definition"</code> (YAML). the properties were copied from <code>Scalable</code> and did not reflect the properties listed in the table. • C.3.8.1: Added <code>"TOSCA version"</code> to OS capability instead of string. TOSCA-134 fix. • A.6.1: Removed TOSCA-225 comments (cardinality) as we now have <code>"occurrences"</code> keyword for this. • F.1.1: Fixed use case to use ONLY the current requirement and capability grammar, removed use cases that showed grammar that was no longer supported. Adjusted text, notes and best practice text to reflect the current design of the normative <code>WebServer</code> and <code>WebApplication</code> types. • A.7.1: Reworked the Requirement assignment grammar and use cases to reflect the desired requirement grammar within Node Templates. Removed 2 examples in this section that are no longer valid and fixed the remaining three examples to be accurate to the new grammar. Note, some additional work may need to be done on the examples, but the grammar SHOULD be final barring any unforeseen problems during TC review.
WD05, Rev10	2015-03-20	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • A.2.6: Added <code>scalar-unit.frequency</code> to allow for properties that provide values that measure in units per second such as CPU clock rate / processing speed (e.g., operations per second). • C.7.2, C.3.3: Moved <code>num_cpus</code>, <code>mem_size</code> and <code>disk_size</code> properties from

			<p>Compute node to Container capability.</p> <ul style="list-style-type: none"> o Added “required: false” to these properties as per WG agreement. • C.3.3: Added “cpu_frequency” property to Container capability. <ul style="list-style-type: none"> o set default .1 GHz, please review. • C.2.4: fix PortSpec. Not supposed to be a list of source and target. • D.1.1: Added non-normative tosca.capabilities.Docker with port mapping support. • G.1: Added ELK use case to overview table
WD05, Rev11	2015-03-31	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • A.2.2.1: “fix_version” portion of the TOSCA version string is now optional. • A.2.6: Added requirement that it is an error if scalar-units do not both have the scalar and unit portion in a declaration. General cleanup to assure we use the prescriptive language for requirement (e.g., SHALL). • A.5.4: Added “additional requirement” requiring search order for capabilities on a target filter to assume a symbolic name first and a type name second to avoid namespace collisions (although collisions should not occur). • A.5.4.5: Clarified “node_filter” example. • A.3.1: Added Network name aliases (primary used in Endpoints). • Appendix E: Inserted chapter that describes minor changes to the CSAR file format over the version described in TOSCA v1.0. Thanks Thomas for authoring. • Ch 7, 14.3, 15, 16, 17.3, 17.4 – Incorporated fixes and address comments received from Chris. • A.2.3: UNBOUNDED definition updated to be clearer. (addressed comment from Chris) • A.5.8 Attribute Defn. – was missing entry_schema (Chris found) • A.7.2: adjusted Requirement Def. description to be correct. • A.7.3: copy had wrong declared type, changed to string. Adjusted ‘copy’ grammar to be single line. • C.3.10.2: BindsTo derived_from updated. • C.5.4: missing Credential property from ConnectsTo definition (grammar). • C.5.5: Changed “AttachTo” type to “AttachesTo” to match all other “active” names used in Relationship Types. This caused cascading changes to many examples and use cases. (Chris) • F.5.2: Port Node Type now needs “relationship” type in definition. (Chris caught) • F.5.3: Linkable needs to derive from Node capability type. (Chris) • G.1.2: Missing topology_template in example (Chris) • G.1.3: Missing “GB” on scalar unit type (Chris)
WD05, Rev12	2015-04-08	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • A.3.1: Network name alias: defined built-on values for application authors to use to describe networks in Endpoints relative to logical public and private networks (regardless of actual name of the network) assigned at runtime. • A.7.3, A.7.3: Requirement assignment and Node Template keyname “target_filter” renamed to “node_filter” to be more generic and reflect exactly the type of TOSCA entity (node) that the filter would be used to select. • B.3.1: token: Added the “token” intrinsic function to extract substrings that are inside a larger string and separated by tokens. • A.5.9: Attribute assignment: Fixed grammar to include the multi-line form that allows a “description” and “value” keyname. • A.8.1.4: The “output” grammar was based upon the property definition which is incorrect; it should be based upon attribute assignment.

			<ul style="list-style-type: none"> • A.5.6.5: Property definition: Moved requirement language buried in the keynames table to the actual “Additional requirements” section. • A.5.11: Operation definition: Added the ability to have both a “primary” artifact (script) and other dependent artifacts for a single operation. • A.5.6: Repository definition: added definition to allow definition of external (artifact) repositories that can be used to host/hold TOSCA artifacts. • A.5.5: Artifact definition: added “repository” keyname to definition to allow artifacts to be located in an external repository. • A.9: Service Template: added “repositories” keyname to allow one or more repository definitions. • D.3.6: <code>tosca.nodes.Docker.Container</code>: a non-normative “Docker friendly” Node Type for use when modeling Docker “containers” which in TOSCA are really “Containees”.
WD05, Rev13	2015-04-28	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • A.7.2: Requirement assignment: Added “properties” keyname to allow property assignments to the relationship keyname. • C.3.4: Endpoint: made “PRIVATE” default for “network_name” for implementations to default to the first IP address found on the first private network if user does not specify otherwise (as an input or in the template). • A.3.3: Added more complete definitions for PUBLIC versus PRIVATE networks. • C.7.10, C.7.11: Container.Runtime and Container.Application replace Container and Containee to better adapt to semantic of the Linux container communities such as Docker and Rocket. Note this mirrors WebServer and WebApplication and can be applied to Java runtimes and apps as well. • H.1.10.3, H.1.10.4: Provided latest “master” service template for ELK and removed scripts as there are too many to reasonable list. Instead we are moving to providing a URL to the live repo. • H1: Completely revamped section to reflect actual running use cases with links to current code. Removed use cases that did not make sense, added others that showed interesting TOSCA concepts. • H1: Added URL links to all use cases master/entry definition Service Template as it is stored in OpenStack for reference. • H.3.2-H.3.6 now lists each Block storage variation (use case) as its own top-level use case. • H.3.8: ObjectStorage use case: Moved up and reworked to new style and latest service template • H.3.9-H.3.11: Network use cases: we now have 3 different variant network use cases thanks to Simeon Monov. • H.1: added Docker use case description • H2: removed outdated Blockstorage use cases which are now covered by new ones or did not make sense (e.g., multiple block storage example). • F.5.1. Network Node: Added network type and physical network optional properties per Simeon Monov suggestion. • C.8.3, C.8.4: Added base artifact types for both implementation and deployment types. • C.7.2: Compute: Changed “Endpoint” to “Endpoint.Admin” • C.7.7: Database: Simplified description of Database node type. • C.8: Added “Type URI”, “Qualified” and “Shortname” names for TOSCA artifact types. • C.8.3: Added Python implementation type. • C.2.1: Credential: No longer has “network” in type URI • C.5.4: ConnectsTo: Use full TypeURI for Credential in definition.

			<ul style="list-style-type: none"> • C.7.12, C.7.13: Added FloatingAddress and LoadBalancer • C.5.6: Added RoutesTo relationship type for FloatingAddress.
WD05, Rev14	2015-04-29	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • B.8.1: Updated get_artifact to include work group's agreed upon parms, Thanks Luc Boutier for this contribution (TOSCA-220). • A.6.2: Requirement Definition: Added support in grammar to add additional Property definitions on known interfaces or their operations. • A.5.11, A.5.12: Updated Interface and Operation definitions to include consideration of when they are used as part of a Requirement definition or Requirement assignment (in a Node Type or Node Template respectively). • A.6.1: Added "occurrences" keyname to Capability def. to support NFV use cases. • H.2.7: ObjectStorage: Added template and diagram • C.3.5: Added Endpoint.Public with experimental support for Floating IP capabilities via new properties. • C.7.12: Added LoadBalancer node type which has an Endpoint.Public which can advantage the Floating IP support. • C.5.6: Added RoutesTo for LoadBalancers primarily (but could be used for any intentional network traversal. • C.3.6: Added Add. Req. to Admin endpoint to required security enforcement where possible
WD05, Rev15	2015-05-04	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • Appendix H: Fixed all examples to use scalar-unit.size for all storage_size inputs. (Chris) F • Appendix H: Fixed "KB" to "kB" even though this is allowed by scalar-unit defintions. (Chris) • Appendix H: Assured that the requirement used for Compute for storage attachment was indeed "local_storage". (Chris) • Appendix G.1.4: Used real Compute "local_storage" requirement, provided a location property value for custom AttachTo type and also used the actual Configure interface and one of its pre_configure operations as suggested (Chris). • C.3.5: Endpoint.Public: clarified descriptions and added additional requirement for DNS name. • C.2.1: Credential: Additional requirements added for token validate. Also, keys and protocol are no longer required properties. • C.2.1: Credential: Added optional user property for legacy use cases • C.2.1.5: Credential: Added new example to show use of Credential to pass in a simple user name / password without using BasicAuth. • C.7.6: DBMS: root_password property is not required since MongoDB and other Database services do not need this; if it is required, such as for MySQL then the subtype can alter it to be required. • C.7.7: Database: user and password properties are no longer required as some databases do not use this concept. • D.3.2: MySQL: root_password is required for MySQL database configuration. • H.2.8: ObjectStorage: Fixed use case template to current spec.
WD05, Rev16	2015-05-06	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • C.3.8.2: File: Now derives from tosca.artifacts.Implementation (type name changed). • H.1.16: Updated to spec. and comments added for future changes that are needed. • D.1.2: Image.VM: Added artifact type for VM images. • D.3.2: MySQL: requires root_passwprd • H.1.18: Container example: Fixed artifact and node types in example's template.

			<ul style="list-style-type: none"> • A.8.1.15: Fixed “groups” grammar. • A.9.1: TOSCA-232: metadata keyname added and several optional keynames of the service template were moved under it. This was a requirement that came from the NFV work group. • C.7.13: LoadBalancer: added occurrences and explained why it allows zero applications on its application requirement.
WD05, Rev17	2015-05-14	Matt Rutkowski, IBM	<ul style="list-style-type: none"> • Note: These changes reflect changes agreed to during the 5/7/2016 half-day Simple Profile spec. review. • A.3.2: Directives: Preserved the keyname “directives” for node templates, but here removed values “selectable” and “substitutable” since we no longer need them as all nodes are “selectable” and “substitutable” was determined to be equivalent to selectable. • Ch. 11: Re-authored to define “concrete” vs. “abstract” nodes and all nodes that are “abstract” in TOSCA are selectable (and therefore substitutable). • Ch. 12: Moved this to Ch. 11.3 as another use case for using node-filter to fulfill an abstract node through selection. • Ch. 13: Substitution for chaining: Removed the need to use the “substitutable” directive” (only use of it anywhere in the spec.); all abstract nodes can be “substitutable” including the Transaction and Queueing subsystems described in the use case in this chapter. • Ch. 3-6: Added logical diagrams to help visualize use cases described only as YAML templates. • D.1.2.1: VM: Added a note that future subclasses of the VM artifact type might include popular standard image and container formats • H.1.16, H.1.17, H.1.18: Updated use cases YAML to latest draft.

