

Configuration Description, Deployment, and Lifecycle Management

Component Model Version 1.0

Status of this Memo

This document provides information to the community regarding the specification of the Configuration Description, Deployment, and Lifecycle Management (CDDLML) Component Model. Distribution of this document is unlimited.

Abstract

Successful realization of the Grid vision of a broadly applicable and adopted framework for distributed system integration, virtualization, and management requires the support for configuring Grid services, their deployment, and managing their lifecycle. A major part of this framework is a language in which to describe the components and systems that are required. This document, produced by the CDDLML working group within the Global Grid Forum (GGF), provides a definition of the CDDLML component model and the process whereby a Grid Resource is configured, instantiated, and destroyed.

CDDLML Specs

There are five documents created by the CDDLML group. They are described in the text below.

The CDDLML Foundation document sets the stage for the remaining documents by introducing the area, by describing functional requirements, use cases, and high-level architecture. It also compares this group with other working groups as well as with other related efforts in the industry.

The SmartFrog Language spec describes a language primarily intended for configuration description and deployment. It is declarative, i.e. it supports attribute value pairs. Furthermore, it supports inheritance, references (including lazy), parameterization, predicates and schemas. It has the same functionality as CDL (see next paragraph), except that it is not XML-based. SmartFrog language predates CDL and it was used as a model when creating CDL. The two languages will be compatible. CDL is primarily intended for machines, SmartFrog for humans.

The CDDLML Configuration Description Language (CDL) is an XML-based language for declarative description of system configuration that consists of components (deployment objects) defined in the CDDLML Component Model. The Deployment API uses a deployment descriptor in CDL in order to manage deployment lifecycle of systems. The language provides ways to

describe properties (names, values, and types) of components including value references so that data can be assigned dynamically with preserving specified data dependencies. A system is described as a hierarchical structure of components. The language also provides prototype-based template functionality (i.e., prototype references) so that the user can describe a system by referring to component descriptions given by component providers.

The CDDLM Component Model outlines the requirements for creating a deployment object responsible for the lifecycle of a deployed resource. Each deployment object is defined using the CDL language and mapped to its implementation. The deployment object provides a WS-ResourceFramework (WSRF) compliant "Component Endpoint" for lifecycle operations on the managed resource. The model also defines the rules for managing the interaction of objects with the CDDLM Deployment API in order to provide an aggregate, controllable lifecycle and the operations which enable this process.

The deployment API is the WSRF-based SOAP API for deploying applications to one or more target computers. Every set of computers to which systems can be deployed hosts one or more "Portal Endpoints", WSRF resources which provide a means to create new "System Endpoints". A System Endpoint represents a deployed system. The caller can upload files to it, and then submit a deployment descriptor for deployment. A System Endpoint is effectively a component in terms of the Component Model specification -it implements the properties and operations defined in that document. It also adds the ability to resolve references within the deployed system, enabling remote callers to examine the state of components with it.



GLOBAL GRID FORUM
office@ggf.org
www.ggf.org

Full Copyright Notice

Copyright © Global Grid Forum (2004). All Rights Reserved.

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 paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the GGF or other organizations, except as needed for the purpose of developing Grid Recommendations in which case the procedures for copyrights defined in the GGF Document process 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 the GGF or its successors or assigns.

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

Intellectual Property Statement

The GGF 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. 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 specification can be obtained from the GGF Secretariat.

The GGF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this recommendation. Please address the information to the GGF Executive Director (see contact information at GGF website).

Table of Contents

Table of Contents	4
1 Introduction	6
1.1 Notational Conventions.....	6
1.2 Namespaces	7
1.3 Definitions.....	7
2 CDDLW-G and the Purpose of this Document	8
2.1.1 Configuration Description Language.....	8
2.1.2 Component Model.....	8
2.1.3 Deployment API	9
3 Components in the CDDLW Architecture	9
3.1 Core Concepts and Terminology.....	10
3.1.1 Components	10
3.1.2 Deployable Systems.....	10
3.2 Deploying Components.....	12
3.3 Component Lifecycle.....	13
3.3.1 State and State Extension Representation	15
4 Components.....	16
4.1 Deployment Components.....	16
4.1.1 ComponentReference	17
4.1.2 CodeBase	17
4.1.3 CommandPath.....	18
4.2 Component Delegates	18
4.3 Systems	19
4.4 Component Properties and Operations	20
4.4.1 Component Properties.....	21
4.4.2 Basic Component Operations	24
4.5 Faults.....	26
4.5.1 Fault Categories	26
4.5.2 Fault Security.....	27
4.5.3 Internationalization	27
4.5.4 DeploymentFault.....	27
4.5.5 Component Faults	28
4.6 Events.....	28
4.6.1 Event Interfaces	29
4.6.2 Event Format	29
4.6.3 Event Registration.....	30
4.6.4 Property Change Notifications	32
4.6.5 Fault Notifications.....	32
4.7 Control Flow.....	32
4.7.1 Sequence.....	33
4.7.2 Reverse	34
4.7.3 Flow.....	34
4.7.4 Wait.....	34
4.7.5 Switch.....	34
4.8 Component-specific message sets and messages	35

5	Reference Resolution	35
6	Using the Component Model	36
6.1	Scope	36
6.1.1	Delegates	36
6.1.2	Document Structure	37
6.1.3	Flow Control Example	38
6.2	Synchronization	39
6.2.1	Events	39
6.2.2	Synchronization Points	40
6.2.3	Events and Synchronization Points	41
6.3	Aggregate System State	41
6.4	Controlling Deployments	42
6.4.1	Event Handlers	43
7	Security Considerations	43
8	Editor Information	43
9	Acknowledgements	44
	References	45
	Appendix A – Component Object Definitions	46
A.1	XML Schema	46
A.2	WSDL 1.1	53

1 Introduction

Deploying a complex, distributed service presents many challenges related to service configuration and management. These range from how to describe the precise, desired configuration of the service, to how we automatically and repeatably deploy, manage and then remove the service. This document addresses the description challenges, while other challenges are addressed by the follow-up documents. Description challenges include how to represent the full range of service and resource elements, how to support service "templates", service composition, correctness checking, and so on. Addressing these challenges is highly relevant to Grid computing at a number of levels, including configuring and deploying individual Grid Services, as well as composite systems made up of many co-operating Grid Services.

1.1 Notational Conventions

The keywords "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119].

This specification uses a notational convention, referred to as "pseudo-schemas" in a fashion similar to the WSDL 2.0 Part 1 specification [WSDL 2.0]. A pseudo-schema uses a BNF-style convention to describe attributes and elements:

- '?' denotes optionality (i.e. zero or one occurrences),
- '*' denotes zero or more occurrences,
- '+' denotes one or more occurrences,
- '[' and ']' are used to form groups,
- '|' represents choice.
- {any} is a placeholder for elements from some other namespace (like ##other in the XML Schema).
- An ellipsis, or three consecutive periods, ..., are used in XML start elements to indicate that attributes from some other namespace are allowed.

Attributes are conventionally assigned a value which corresponds to their type, as defined in the normative schema.

```
<!-- sample pseudo-schema -->
<element
  required_attribute_of_type_QName="xsd:QName"
  optional_attribute_of_type_string="xsd:string"? >
  <required_element />
  <optional_element />?
  <one_or_more_of_these_elements />+
  [ <choice_1 /> | <choice_2 /> ]*
</element>
```

A full XML Schema and WSDL description of all elements and operations are available in Appendix A of this specification. Sections 4 and 5 of this specification and appendices

A.1 through A.4 containing the XML Schema (XS) of the component model and the WSDL component interface are *normative* specifications. All UML diagrams found are non-normative, and meant to serve only as illustrations of the CDDLM concepts.

1.2 Namespaces

The following namespaces are used in this document:

xsd	http://www.w3.org/2001/XMLSchema
wsa	http://schemas.xmlsoap.org/ws/2003/03/addressing
wsbf	http://docs.oasis-open.org/wsrfl/2004/06/wsrfl-WS-BaseFaults-1.2-draft-01.xsd
wsrp	http://docs.oasis-open.org/wsrfl/2004/06/wsrfl-WS-ResourceProperties-1.2-draft-01.xsd
wsrl	http://docs.oasis-open.org/wsrfl/2004/06/wsrfl-WS-ResourceLifetime-1.2-draft-01.xsd
wsnt	http://docs.oasis-open.org/wsrfl/2004/06/wsn-WS-BaseNotification-1.2-draft-01.xsd
wstop	http://docs.oasis-open.org/wsrfl/2004/06/wsn-WS-Topics-1.2-draft-01.xsd
muws-p1-xs	http://docs.oasis-open.org/wsdm/2004/12/muws/wsdm-muws-part1.xsd
muws-p2-xs	http://docs.oasis-open.org/wsdm/2004/12/muws/wsdm-muws-part2.xsd
mows-xs	http://docs.oasis-open.org/wsdm/2004/12/mows/wsdm-mows.xsd
cdl	http://www.gridforum.org/namespaces/2005/02/cddlm/CDL-1.0
cmp	http://www.gridforum.org/cddlm/components/2005/02
cmpw	http://www.gridforum.org/cddlm/components/2005/02/wsd
api	http://www.gridforum.org/cddlm/deployapi/2005/02

1.3 Definitions

The following terms are defined and used throughout this document. This section is intended as a brief reference only of these concepts.

Resource – A WS-Resource Framework construct to define an identity, properties and standard interaction through a Web Service to access a logical object.

Service – A WS-I compliant Web Service

Service Endpoint – The WSDL defined interface of an instantiated Web Service.

Component – The smallest deployment unit within the CDDLM framework.

Language Component – A deployment construct defined within a CDL document which does not correlate to any underlying object or resource during deployment.

Deployment Component – A deployable resource defined within a CDL document which will be instantiated during deployment.

Component Delegate – A deployment component which subsumes the responsibility of the lifecycle of child components from the rest of the Framework.

System – A group of Components defined and named in a CDL deployment document.

Deployment Graph – An internal representation of the relationships among Components.

Portal EPR – The WS-Resource of the Deployment API portal interface.

System EPR – The WS-Resource of the node selected by the Deployment API to act as the root of a deployment.

Component EPR – The WS-Resource of a specific Component within a deployment.

2 CDDLML-WG and the Purpose of this Document

The CDDLML WG addresses how to: describe configuration of services; deploy them on the Grid; and manage their deployment lifecycle (instantiate, initiate, start, stop, restart, etc.). The intent of the WG is to gather researchers, developers, practitioners, and theoreticians in the areas of services and application configuration, deployment, and deployment life-cycle management and to explore the community need for a broader effort in this area. The target of the CDDLML WG is to come up with the specifications for CDDLML a) language, b) component model, and c) deployment API. This document represents the component model specification. This specification is derived from the SmartFrog framework developed at HP Labs, and other industry specifications.

This document describes the component model of the framework. It does not describe any required component libraries. The design of the system is loose to allow the binding of attributes to deployment objects to be handled by an individual implementation of the system. This document serves to outline the requirements for a software object to be deployable by the CDDLML framework.

This specification is closely related to a set of three specifications that together comprise service configuration description, deployment, and lifecycle management. All three specifications are briefly described in the following three paragraphs. For more details please refer to other two specifications.

2.1.1 Configuration Description Language

The CDDLML Configuration Description Language (CDL) [**XML-CDL**] is an XML-based language for declarative description of system configuration that consists of components (deployment objects) defined in the CDDLML Component Model. The Deployment API uses a deployment descriptor in CDL in order to manage deployment lifecycle of systems. The language provides ways to describe properties (names, values, and types) of components including value references so that data can be assigned dynamically with preserving specified data dependencies. A system is described as a hierarchical structure of components. The language also provides prototype-based template functionality (i.e., prototype references) so that the user can describe a system by referring to component descriptions given by component providers.

2.1.2 Component Model

The CDDLML Component Model outlines the requirements for creating a deployment object responsible for the lifecycle of a deployed resource. Each deployment object is defined using the CDL language and mapped to its implementation. The deployment object provides a WS-ResourceFramework (WSRF) compliant "Component Endpoint" for lifecycle operations on the managed resource. The model also defines the rules for managing the interaction of objects with the CDDLML Deployment API in order to provide an aggregate, controllable lifecycle and the operations which enable this process.

2.1.3 Deployment API

The Deployment API [**CDDL**M-API] is the WSRF-based SOAP API for deploying applications to one or more target computers. Every set of computers to which systems can be deployed hosts one or more "Portal Endpoints", WSRF resources which provide a means to create new "System Endpoints". A System Endpoint represents a deployed system. The caller can upload files to it, then submit a deployment descriptor for deployment. A System Endpoint is effectively a component in terms of the Component Model specification - it implements the properties and operations defined by the model. It also adds the ability to resolve references within the deployed system, enabling remote callers to examine the state of components with it.

3 Components in the CDDL M Architecture

A prerequisite for understanding this document is to have read the CDDL M language document(s), which introduce a tree-like representation of configurable components as a way of describing the pieces of a distributed application. Using the description of the application, a system definition can be provided to an instance of the CDDL M Deployment API . It is the responsibility of this instance to realize the deployment through instantiation and control of CDDL M compliant deployment components. This basic relationship can be understood from this simple figure:

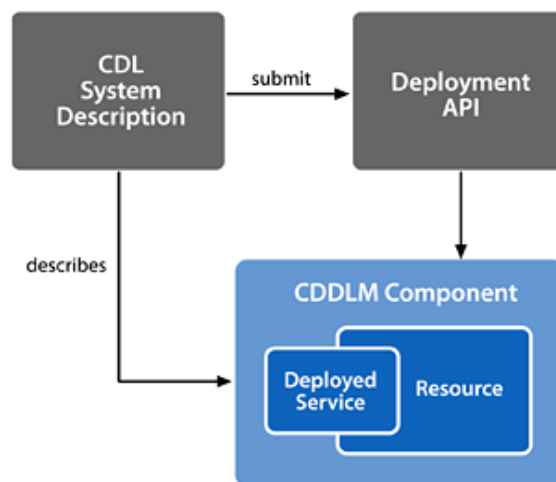


Figure 1. The CDDL M Deployment Framework.

The CDDL M component model is not intended to be a general purpose model for systems or application management. It is limited in scope to address the problems of deploying and managing the deployed state of *resources*. In the context of CDDL M, a resource is considered to be the hardware, software and other configurable items required to execute the distributed application being managed.

The basic principles of the CDDLM framework are described in the CDDLM Foundation Document. Additionally, the reader should consult the CDDLM Deployment API document in order to understand the basic deployment capabilities of the CDDLM framework.

3.1 Core Concepts and Terminology

3.1.1 Components

The CDDLM component model represents a means to group a hierarchy of services within a deployment container. This container provides a wrapper for the services to aid in their deployment and management. In this containment hierarchy there may be a many-to-many relationship among services and actual instantiated resources. This hierarchy is not intended to reflect the organization of those managed resources. Its purpose is to capture the runtime deployment dependencies, operational flow and state of those resources.

At deployment time, a service will be resolved by the CDDLM implementation to determine its composition and membership within the active Grid Fabric. The basic unit of deployment is a *component*. The service to be deployed and its resource will be a part of a component which bounds its contents and defines its deployment behavior.

A component has two analogues. A *language component* is a configurable component along with its property list as defined in the CDL language document; it is a declaration of what the user wishes deployed, and how they want it configured. A *deployment component* is a code library that provides the implementation of language components with interfaces and operations defined later in this document. When the CDDLM infrastructure needs to deploy a service, it must locate and instantiate deployment components for every language component that it wishes to instantiate. These deployment components will be moved through their lifecycle and so bring up the service.

A CDDLM *component object* is a specific instance of a deployment component. A deployment component can be instantiated multiple times in the same document or system, as well as in multiple, separate systems. The component object will correlate to a component within the system tied to one or more allocated resources. A deployment component will have one component object for each resource or group of resources that it is declared to manage.

3.1.2 Deployable Systems

Groups of components can be organized to form a deployable *system*. A system will be comprised of one or more components each implementing some portion of a complete Grid Service or prerequisite, or individually implementing one or more *services*. Each service will provide access to its implementation through one or more *endpoints*, either WSRF endpoints or standard WS-I endpoints, as defined in the services own WSDL. All such endpoints will be presumed to be valid Endpoint References (EPRs) as defined by

the WS-Addressing specification. The relationships between services and components are modeled in the language and exposed in order to facilitate these compositions.




Figure 2. A Deployable System.

Not all services will be complete software applications. A service endpoint may exist within a larger set of an application server, for example. A service may also be dependent upon the configuration of a more fundamental portion of software such as an operating system or patch. Or there may be a pre-existing service that does not implement the CDDLM interfaces. In this manner, the CDDLM component model enables services to be “wrapped” within a component container that is responsible for implementation of the defined component interfaces.

In many cases, the services that are deployed do not meet the manageability requirements of its environment. The CDDLM component model allows a component to wrap an implemented service and provide a level of manageability that otherwise would not exist. These manageability functions implement CDDLM specific interfaces as well as those required by the WSDM specifications. Other interfaces could be provided as the implementation dictates.

Though the conceptual diagram above depicts a system as shown hierarchically, there is no requirement by CDDLM for components to be explicitly organized in this way. A system is just a loose confederation of components needed to deploy some set of endpoints. The following UML diagrams the conceptual relationship of these entities.

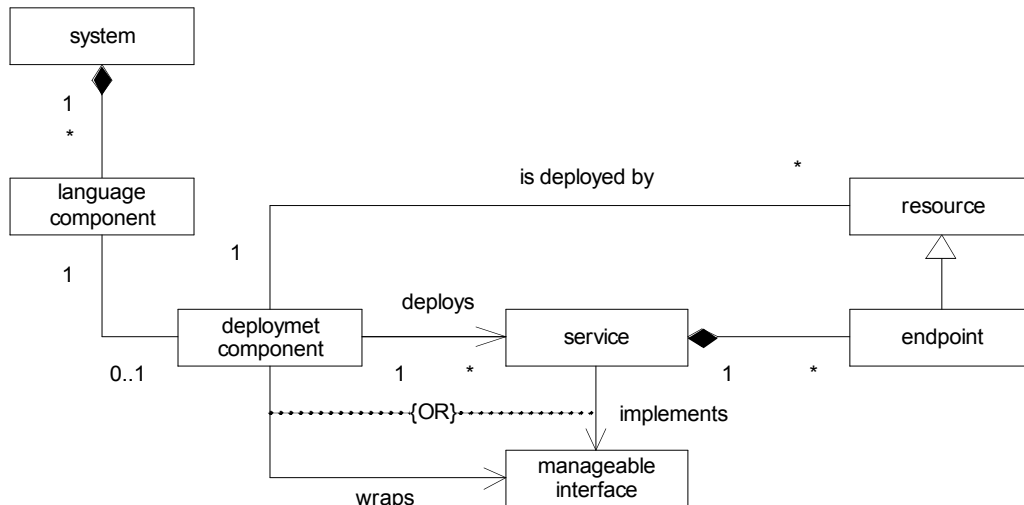


Figure 3. Component Model Conceptual Diagram.

3.2 Deploying Components

In order to deploy, the CDDL framework is sent a description of the system in a supported language through a node implementing the Deployment API. This node is addressed via a well known WS-Addressing Endpoint Reference (EPR) [WS-A], a *portal EPR*. The run time of this node parses the description in preparation to begin the deployment. As defined in the CDL Language Document, the description will always be translated to XML-CDL during parsing.

The output of the parsing stage is the production of an internal representation of the system. This document will use a (annotated) tree of language components -a *deployment graph* - as an exemplary structure. The Deployment API does not require that this graph be created or maintained. We use this formalism here to describe how the deployment will interact both within components and in conjunction with the Deployment API.

This graph describes what must be deployed, and what the interdependencies are. Using the relationships defined in the CDL document, the basic hierarchy that is implicit in the structure of the document can be transformed into a graph showing the explicit relationships, ordering and timing of deployment operations and events. This graph can then be used, conceptually, by the Deployment API to create and enforce these relationships at deployment-time.

Each system to be deployed is instantiated through the creation of a WS-Addressing Endpoint Reference to a resource responsible for its operation. This *system EPR* represents the aggregate state of the collection of deployed components for a particular system instance. This EPR provides a WS-ResourceFramework resource with all of its methods and properties exposed using the WSRF embodiment of the service endpoint. The complete description of using the *system EPR* is described in the CDDL Deployment API.

The resource responsible for the *system EPR* or *portal EPR* will parse the CDL definition to map language components to deployment components. The parsing phase produces a complete XML CDL representation of the deployment graph. All internal use of documents and languages by the component model are based on the XML CDL specification, no other languages are supported. The deployment graph will be segmented and for each deployment component identified a WS-Addressing endpoint reference, a *component EPR*, is assigned to each deployment component.

A *component EPR* is the primary interface to manage a deployment component. Component EPRs are references to the resource which is an instance of the corresponding deployment component. This EPR will also rely on the mechanisms defined by the WS-ResourceFramework. There is no requirement of this component model or by CDDLM that a component map to only a single resource, or that the address of the component endpoint correlate in any way to the deployment location of the resource. Thus, the embodiment of “WS-Addressing Using Reference Properties” MUST be followed. All message exchanges with a component will include the ReferenceProperties of the WS-Resource qualified component EPR.

For each deployment component, the framework must then instantiate each component object. During this instantiation, the component must become able to respond to operations at its endpoint. The endpoint, once available, is used to transfer the fragment of CDL relevant to the component object and initialize base properties of the component. The component must also be notified of the endpoint address of the system EPR for later use.

Once deployed, the component EPRs are used internally for components to communicate with the system EPR and with each other. The node responsible for the system EPR is able to send lifecycle commands to component EPRs and receive lifecycle events. The system EPR provides endpoint resolution services to components in order to enable this intercommunication. This resolution service allows components to translate CDL language references to another component into the EPR of that component. This resolution interface is described in detail in the Deployment API specification.

3.3 Component Lifecycle

The basic lifecycle of components, and thus the resources being managed, is defined by a state machine that is implemented by all component objects within CDDLM. The CDDLM lifecycle is restricted, in that it supports the deployment processes used by the framework and models only the deployment state of the system, not its operational characteristics. For example, once a CDDLM component has entered the running state, CDDLM considers it to be available. In reality, the service that is being deployed can be operative but have suffered some set of internal failures.

Each deployment component independently represents the state of the deployed resource which it is managing. The system as a whole must also represent a reasonable depiction of the overall state of many components. In this state machine, there is no representation

of transitional state, only of arrival at new states. The core lifecycle is defined by the states, allowed transitions and operations shown in the diagram in below.

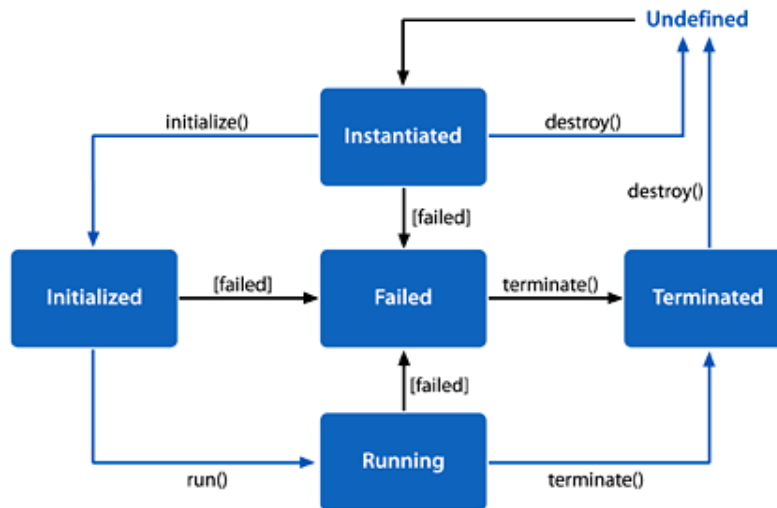


Figure 4. CDDLM Component Lifecycle Model.

The *undefined* state represents any state in which the component is not known to have an instance running or instantiated in the system. In the undefined state, the assets of the component may or may not have been deployed from the rest of the framework. There is no requirement by the component model that the assets of the system be stored or removed during any lifecycle operation.

When a component is created, it will make the transition from undefined to *instantiated*. During this time, the component and its deployment package will be deployed to the appropriate system, the component instantiated and any operations will be performed that are part of the component's instantiation process. This state also presumes that whatever activation is required in order for the WSRF resource identifier of the component to be valid has been performed. Operations on the endpoint reference can now be performed. As shown in the diagram above, the `initialize()` and `destroy()` state change commands are supported in this state.

The next state of the system is the *running* state. This state indicates the services provided by the resources that are being deployed are available for use. As mentioned earlier, this state does not indicate any information regarding the operational capabilities of the service, only that it has completed initialization and not failed.

At any time, state actions may not complete correctly or the service itself may fail. In response to these failures, the component will transition to the *failed* state. The component may remain active in the system, but its managed resource is presumed to no longer be operational.

Once a component is running or has failed, it must be terminated to stop its services. The *terminated* state represents a state where a component is no longer running and cannot be

returned to the running state without redeployment. This state, however, does not eliminate the resource from the system. Upon invocation of the `destroy()` command, the component will return to the undefined state and the corresponding resource becomes invalid.

In a system with multiple components, the lifecycle of the whole system is defined by the relationships between the individual component lifecycles. The hierarchy of the system defines relationships where related components' lifecycles are linked. The component model and the CDL language help define explicit semantics for guiding lifecycle transitions. These transitions are defined in detail in Section 5.3 of this document.

3.3.1 State and State Extension Representation

It is also desired for the lifecycle model to be extensible at the component level. The core lifecycle described above is required for implementation of CDDL. However, CDDL will enable services to declare service specific lifecycle extensions that are declared for a particular component and operate within the basic deployment lifecycle model defined herein. In this way, a service can be deployed and managed simply, but also define internal states relevant to other services capable of consuming its extended information.



Figure 5. Extended States.

In the example above, a service is currently known to CDDL and has been created and is running. If it is sent a command such as “Refresh Data”, it may externally remain in the running state. However, internally it may go through several state transitions in order to execute the command. By declaring these states in a resource property document, its states can be tracked and understood through the basic component model interfaces.

This extension mechanism does not require the definition of the state machine or state transitions to the CDDL framework. This would require implementation of a formal calculus within the CDL language, and a secondary state machine within the framework. Instead, the lifecycle model is implemented in the component model using the WSDM MUWS State capability. This allows the state of the system to be exposed as a WSRF resource property with a simple and flexible model for extension.

4 Components

This section describes the model for deployment components within the CDDL framework: what they are, how they work, and how they integrate with the framework. The CDDL framework requires a concrete object that can be instantiated in order to perform a deployment action. This object's purpose is to abstract the basic framework from the details of the composition of what is to be deployed and its specific, required deployment steps.

Components are represented in the CDL language using an abstract object model that most closely resembles languages such as Scheme: any component can have new attributes added to it dynamically, while the type system is very loose. Furthermore, there is no split between classes and instances; everything is an instance that can be extended dynamically.

There also does not exist a mapping between every component instance in the language-level model and implementation model. That would be neither sensible nor possible. Instead, any component in the language can be entirely abstract. This has benefit for providing templates and parameterization. Language components can also be used to provide higher level behaviors to coordinate deployments. For a list of supported language component behaviors, consult the CDL Language Specification Document.

4.1 Deployment Components

In the CDDL system, a CDL document will contain one or more language components. In order to turn this into a deployment, the framework will need to map these language components into deployment components. When a `create()` command is sent to the Deployment API, the system will parse the CDL document and create a mapping.

At least one of the language components declared in the CDL document **MUST** map to a deployment component. Otherwise, no action can be taken during the deployment phases, as all concrete actions of the framework are performed against the deployment components of a system. Any component object **MAY** simply exist to assist with the deployment and have no actual internal objects to deploy or resources to manage. Each component object, though, **MUST** implement the basic interfaces defined in this specification document.

The parsed CDL document can be walked to identify deployment components and assemble a deployment graph and actions to perform in order to instantiate components. For each element that declares an asset location and means to activate the assets, an object will be created. The runtime will deploy the code contained at the source location, and attempt to instantiate the object. An example of how this is done with CDL is shown below.

```
<cdl:cdl targetNamespace="http://example.org/webapp-template">
<cdl:system>
<WebApplication>
  <cmp:CodeBase>http://server/file.jar</cmp:CodeBase>
```



```
<cmp:CommandPath>com.exns.ApacheDeployer</cmp:CommandPath>
<port>80</port>
<hostname>www.example.org</hostname>
</WebApplication>
</cdl:system>
</cdl:cdl>
```

In this case, the `WebApplication` corresponds to an `ApacheDeployer` deployment component located within the deployment container `file.jar` at the remote web server location.

Once deployment components have been identified, the processing is very platform specific. A Java deployment component will be activated within some JVM to become a component object. An executable will be run on its host. A COM or other object will be created within some host specific to the platform. It is important that during the resolution and instantiation process that the CDDLM framework is able to bind the component to the correct runtime or host. Further parameterization of the host requirements should be given in the CDL for the component.

There are a base set of component properties that define a deployment component within a CDL document. Each one of these properties, if declared within the CDL document, MUST be exposed as resource properties of the component as defined by the WS-ResourceProperties specification

4.1.1 ComponentReference

The component endpoint reference (EPR) SHOULD be specified within the CDL document. This element definition enables the CDDLM framework to identify a target resource's endpoint. During the resolution phase of parsing the CDL document, as described in the CDL language document, CDDLM MUST identify the target resources for the deployment and assign them unique endpoint references.

The `ComponentReference` CDL property is shown in the following XML representation. It MUST be a direct child of the language component it modifies. It cannot be multiply defined for a single component.

```
<ConfigurationElement>
  <cmp:ComponentReference>
    wsa:EndpointReferenceType
  </cmp:ComponentReference>
</ConfigurationElement>
```

As the component object is responsible for one or more managed resources, its endpoint reference MUST include the `ReferenceProperties` child element with a stateful resource identifier. This will ensure that the component object resource is clearly identified separately from the resource it is managing.

4.1.2 CodeBase

The component code base is defined as any location of file assets at some URI. It is up to the Deployment API how to obtain and upload the file(s). Any CDL element which is declared as part of the system description, a child of the `<cdl:system>` element, is potentially a deployment component. If such an element contains the `CodeBase` property

as a direct child of the element, that element **MUST** be considered a deployment component. If the same element does not contain a `CommandPath` property to indicate how to activate the component, an error **MAY** result.

The `CodeBase` CDL property is shown in the following XML representation. This property can be defined as many times as needed to indicate all of the possible sources of deployment assets.

```
<ConfigurationElement>
  <cmp:CodeBase>xsd:anyURI</cmp:CodeBase> *
</ConfigurationElement>
```

4.1.3 CommandPath

The command path is some referenceable path indicating to the Deployment API how and/or what to instantiate in order to realize this deployment component. This path is defined as an opaque string as it may be any one of a number of things such as a Java class path, an executable name, a COM object name or GUID, or any other platform specific instantiation name.

The `CommandPath` property can also include sub-elements that can be used as arguments or other augmentation of the base path. This allows the CDL to declare either a structured command using path and arguments, or an unstructured opaque string. Both can also be defined, but if the structured path form is used, the path element is not optional.

The `CommandPath` CDL property is defined as a `cmp:commandPathType` element which is shown in the following XML representation.

```
<ConfigurationElement>
  <cmp:CommandPath>
    {any}
    <cmp:path>xsd:string</cmp:path>
    <cmp:args>...</cmp:args> *
  </cmp:CommandPath>
</ConfigurationElement>
```

4.2 Component Delegates

A derived type of component is a *component delegate*, a deployment component that manages a set of child components on behalf of the system. Component delegates extend the basic functionality of a component by aggregating the lifecycle and operations of components that are declared to be their children. Within the deployment, a component delegate presents a single endpoint and resource for management. However, internally all component children **MUST** still present a full component endpoint for management by the delegate.

In order to create a component delegate, you will simply declare that the component is a delegate. There are no attributes or children of this declaration. However, within the CDL document, a component delegate will need its children defined. Any CDL element which is identified as a deployment component will have a property list defining the parameters of the component. The element may also have child elements as part of the

property lists. Within this hierarchy, if the element is declared as delegate, any child element which is also a deployment component will be a child of the component delegate. Within a CDL document, a declaration may look as follows, with a database server declared as a child of the web server component delegate:

```
<cdl:cdl targetNamespace="http://example.org/webapp-template">
<cdl:system>
<WebApplication>
  <cmp:CodeBase>http://server/file.jar</cmp:CodeBase>
  <cmp:CommandPath>com.exns.ApacheDeployer</cmp:CommandPath>
  <cmp:Delegate/>
  <port>80</port>
  <hostname>www.example.org</hostname>
  <OracleServer cdl:extends="DBServer">
    <cmp:CodeBase>http://server/oracle.zip</cmp:CodeBase>
    <cmp:CommandPath>./oraInst/install.sh</cmp:CommandPath>
    <username>oracle</username>
    <password>sa</password>
    <tnsname>server.exns.com</tnsname>
  </OracleServer>
</WebApplication>
</cdl:system>
</cdl:cdl>
```

If a component is declared to be a component delegate, but has no children, it **MUST** operate as a standard component. If a component has children, but it itself is not a delegate, there is no parent-child or other relationship and the component **MUST NOT** enforce any such dependency.

4.3 Systems

Within a CDL document, all of the elements defined by this specification will be applied as part of normal property lists, as children of CDL configuration or system elements. This is done to keep the deployment process description within the declaration of the deployment components themselves. This will allow us to use the standard parsing features of the CDL language to our benefit. The basic component model as defined thus far and relationships are characterized in the following diagram.

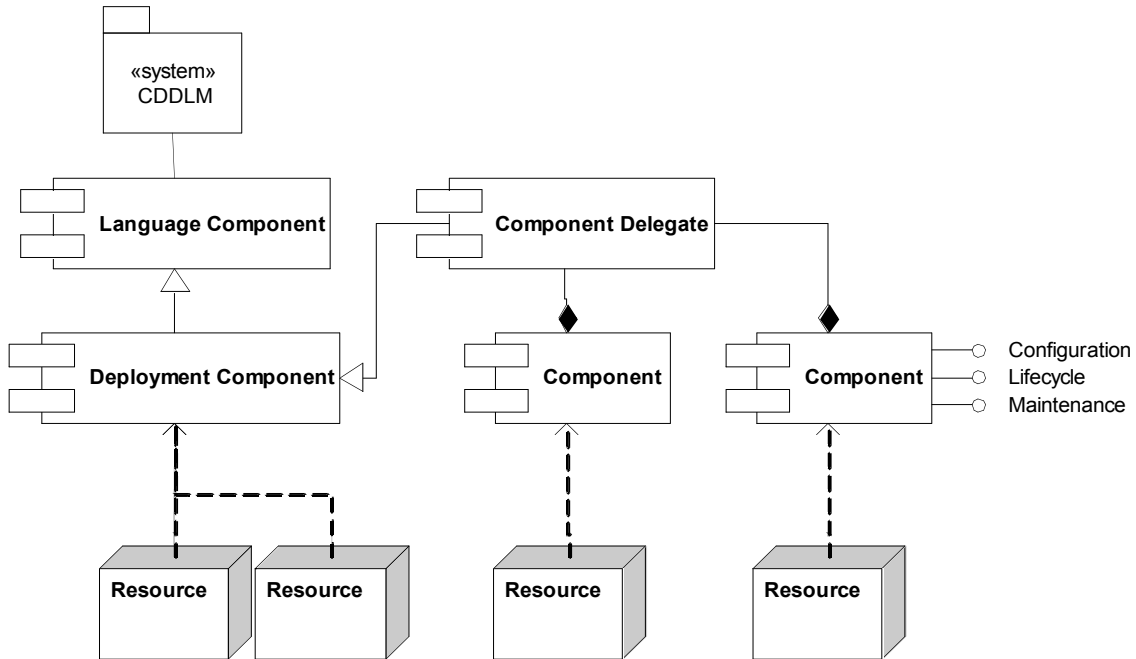


Figure 6. Basic Component Relationships.

The attribute inheritance model of CDL implies that components which derive their configurations from others must be able to override those attributes. Though this implies some derivative behavior in the language, the component model does not require the objects to be in a parent-child relationship or have some inherited behavior in its implementation. Thus there is no requirement for any form of object inheritance or other object oriented implementation.

A system is simply a bag of elements, component objects. There may be a hierarchical, peer or other relationship between objects, if deployment dependencies exist. Without dependencies the CDDL runtime is free to deploy in any means it chooses. Any component object not related to another is a peer of all other top level components. Other than the dependent relationship, there is no other implied constraint between component objects.

4.4 Component Properties and Operations

Every deployment component managed by the CDDL framework will support a common set of properties and operations. There may also be optional sets of messages, for use between components that have explicit knowledge of each other's message support. Callers **SHOULD** be able to poll endpoints to verify that a set of messages is supported before sending one or more messages from that set.

In addition to the core lifecycle functions, the components will be able to use the CDDL framework to access their configuration information, locate other components

defined in the configuration, and potentially alter the running configuration. These interfaces will be implemented using the WSRF and WSDM specifications where appropriate, as well as companion specifications as indicated.

4.4.1 Component Properties

When a component is defined, its CDL property list defines several attributes related to its configuration. The component **MUST** expose each of these attributes as a resource property of the component WS-Resource. In addition to the properties associated with those defined in the CDL property list, every component has properties and actions for managing its lifecycle.

4.4.1.1 Identity

Every component must be uniquely identifiable to enable proper construction of and communication with its WSRF endpoint. The WSDM MUWS defines the Identity capability to meet this goal. A deployment component **MUST** implement this property and declare a unique URI as defined in the WSDM MUWS Part 1 specification. The WSDM Identity capability defines the ResourceId property for each component.

```
<muws-p1-xs:ResourceId>xsd:anyURI</muws-p1-xs:ResourceId>
```

4.4.1.2 ManageabilityCapability

The WSDM MUWS also requires that a compliant endpoint be able to identify the capabilities it exposes for management. A deployment component **MUST** implement the WSDM ManageabilityCharacteristics capability. This is implemented as a list of ManageabilityCapability(s), each of which contains a URI defining the capability supported by the component.

```
<muws-p1-xs:ManageabilityCapability>xsd:anyURI</muws-p1-xs:ManageabilityCapability> *
```

A component **MUST** at least return the Identity, State and ManageabilityCharacteristics facilities for this property as shown in this XML fragment.

```
<muws-p1-xs:ManageabilityCapability>
http://docs.oasis-open.org/wsdm/2004/12/muws/capabilities/Identity
</muws-p1-xs:ManageabilityCapability>
<muws-p1-xs:ManageabilityCapability>
http://docs.oasis-open.org/wsdm/2004/12/muws/capabilities/State
</muws-p1-xs:ManageabilityCapability>
<muws-p1-xs:ManageabilityCapability>
http://docs.oasis-open.org/wsdm/2004/12/muws/capabilities/ManageabilityCharacteristics
</muws-p1-xs:ManageabilityCapability>
```

4.4.1.3 ComponentName

Every component will also have its own unique name which can be created from its own valid *cdl:ref* reference. This path name **SHOULD** be exposed to help uniquely define the component and provide confirmation of the CDDL source in message exchanges. This attribute is defined as shown in this XML representation.

```
<cmp:ComponentName>cdl:pathType</cmp:ComponentName>
```

The component model does not require that this property be set in any particular manner.

4.4.1.4 *ComponentDescription*

When a component is instantiated, the CDL document may define multiple properties in its property list. These properties can contain name/value pairs for attributes of the component or other information such as the control flow elements or events described later in this document. In order to be instantiated properly, the component must know these declarations. Through the Deployment API, an implementer of the framework **MUST** ensure that after a component has been instantiated, its CDL must be set appropriately.

An implementer of this component model and the framework **MAY** choose to pass the entire CDL document or just the relevant fragment. Thus, every component is **RECOMMENDED** to be able to parse a CDL document and extract its corresponding node. If a component object is a delegate, the descriptor fragment that is passed to the component **MUST** include the CDL of the child nodes. The delegate will be responsible for breaking up and/or passing on the CDL to these nodes.

```
<cmp:ComponentDescription>{any}</cmp:ComponentDescription>
```

4.4.1.5 *ComponentStatus*

Each deployment component **MUST** expose a state resource property which implements the MUWS State capability. To satisfy this requirement, a deployment component must contain *muws-p2-xs:State* and *muws-p2-xs:StateTransition* elements. Additionally, a deployment component may include additional information as an opaque quantity that an external consumer may be able to process. The ComponentStatus property will be exposed by every component object of a system.

```
<cmp:ComponentStatus>
  <cmp:State>UndefinedState|InstantiatedState|InitializedState|
    RunningState|FailedState|TerminatedState</cmp:State>
  <cmp:LifecycleTransition>muws-p2-xs:StateTransition</cmp:LifecycleTransition>
  <cmp:ExtendedState>{any}</cmp:ExtendedState> ?
</cmp:ComponentStatus>
```

The State element implements MUWS State by defining the lifecycle states of CDDL as a list of *muws-p2-xs:StateType* elements. These elements are listed in the table below.

Element	Description
UndefinedState	The undefined state.
InstantiatedState	State representing the presence of a component instance.
InitializedState	State in which a component has been properly initialized.
RunningState	Operational state.
FailedState	State in which the component has failed either a lifecycle operation or its operation has failed.
TerminatedState	State in which a component instance has been terminated.

As the failed state may have been arrived at due to failures during many parts of the lifecycle, it is **RECOMMENDED** that the component take action to ensure the services of

the resources are not available while in this state, particularly if the transition occurred from the running state.

The State element follows the MUWS State capability which allows one to create state extensions according to the rules of MUWS Category Taxonomies. These rules **MUST** be followed, particularly that any extension to the state model listed in this document **MUST** be declared in a resource property document with a valid QName to identify the semantics of the extensions. Also, state extensions must be done by declaring a state to **contain** the states declared in this component model. A schema using this capability would define its own sub states as follows this XML Schema fragment.

```
<xsd:element name="ApplyingPolicy">
  <xsd:complexType>
    <xsd:complexContent>
      <xsd:restriction base="cmp:lifecycleStateType">
        <xsd:sequence>
          <xsd:element ref="cmp:RunningState"/>
        </xsd:sequence>
      </xsd:restriction>
    </xsd:complexContent>
  </xsd:complexType>
</xsd:element>
```

As an example, the “Applying Policy” sub-state from the example in Section 3 may be represented by the following XML fragment.

```
<cmp:ComponentStatus>
  <cmp:State>
    <exns:ApplyingPolicy xsi:type="cmp:lifecycleStateType">
      <cmp:RunningState/>
    </exns:ApplyingPolicy>
  </cmp:State>
  ...
</cmp:ComponentStatus>
```

The LifecycleTransition element is defined as a *muws-p2-xs:StateTransition* element. This state transition property **SHOULD** be implemented as defined by WSDM MUWS. The state TransitionIdentifier **SHOULD NOT** be used, but instead the PreviousState element **SHOULD** be used. An example XML fragment of this property could be as follows.

```
<cmp:LifecycleTransition>
  <muws-p2-xs:StateTransition Time="2005-03-01T01:54:30Z">
    <muws-p2-xs:EnteredState>
      <cmp:RunningState/>
    </muws-p2-xs:EnteredState>
    <muws-p2-xs:PreviousState>
      <cmp:InitializedState/>
    </muws-p2-xs:PreviousState>
  </muws-p2-xs:StateTransition>
</cmp:LifecycleTransition>
```

The final element, ExtendedState, allows implementation specific information to be returned as part of the property. A consumer is not required to be able to process this information.

4.4.1.6 *Notification Properties*

The WS-BaseNotification specification requires that a resource that wants to advertise any notifications to potential consumers **MUST** implement three properties.

```
<wsnt:Topic> *  
<wsnt:FixedTopicSet>  
<wsnt:TopicExpressionDialects> *
```

The Topic resource property contains a collection of Topics supported by the component. The set of topics is expressed using one or more wsnt:Topic resource property element(s). The FixedTopicSet resource property indicates if the collection of Topics defined in the Topic property can change or are static over time. The final property, TopicExpressionDialects, indicates the TopicExpression dialects supported by the component for notifications.

4.4.1.7 *Deployment properties*

In order to facilitate discovery operations, a deployment component needs to provide access to its Resource Property document. The basic properties of a component are declared within the ComponentProperties element. A component **MAY** additionally support the WS-MetadataExchange specification to enable dynamic retrieval of a component's property declarations. This element is normatively declared in the XML Schema in Appendix A.

4.4.2 **Basic Component Operations**

In order to access a component's attributes and work with its management functions, a class of operations are exposed to enable these behaviors. These behaviors are expected to be used both internally by the CDDLM framework and with external systems.

4.4.2.1 *WS-Resource Properties*

Component properties in the component model can be directly manipulated using the WS-ResourceProperties specified mechanisms. Each component **MUST** implement the GetResourceProperty operation and **SHOULD** implement GetMultipleResourceProperties to assist in scalable access patterns.

A component or the system EPR host may choose to alter the value of a component's property. This is most likely as the result of some event or decision step in the deployment process. For completeness, the component **SHOULD** support each of the Insert, Update and Delete operations of the WS-Resource.

In order to support enumeration and introspection type operations on the component's attributes, a component **MAY** implement the QueryResourceProperties operation. It is **RECOMMENDED** that an implementation supports the XPath 1.0 query expression dialect at a minimum, additional query expression dialects **MAY** also be supported.

Optionally, a component EPR or the system EPR may provide a means to retrieve the component's current WSDL which can be dynamically generated to reflect the component's current state. That operation is not normatively defined.

Each component object **MUST** support lifecycle state observation. This interface will use the WSRF property query ability to retrieve the `ComponentStatus` property of the component.

4.4.2.2 State transition operations

During the life of an object, the CDDL framework will invoke state transition actions on the object. In order to follow the WSRF recommendations, each object will provide a separate interface for each state transition request: initialize, run, terminate, and destroy. An instantiate message is unnecessary, as the `create()` operation handled by the Deployment API will have instantiated each component object.

Note that the `destroy()` operation is a WS-ResourceLifetime operation, not a CDDL specific implementation. It is **NOT RECOMMENDED** to implement scheduled termination as part of the WSRF implementation for components.

Each of these messages take no parameters in their input, and return no data. If a failure occurs, a fault of type `StateActionFault` is thrown. It can also be expected that a component will transition to the failed state in this event. All of these operations **MUST** be performed synchronously. If CDDL consumers wish to find out more about the state transition, it must issue a state query request.

Upon invocation of a particular state command, a component may take any appropriate action and alter its referenced state. If no action is to be taken, the component **MUST** at least guarantee transfer to the next state. A component **SHOULD NOT** report its state as the next state until the transition is complete. This conservative approach prevents reporting a successful transition followed by a failure due to problems transitioning to the requisite state.

4.4.2.3 Maintenance operations

It is often the case that simple maintenance activities can prolong the healthy lifespan of a long-running service. It is desired to allow CDDL objects to provide a simple command execution infrastructure to enable periodic maintenance or cleanup of these services. This interface will presume that the object is knowledgeable in means to execute the command. This interface does not replace the use of structured operations declared additionally by the component, but serves as a simple, well-known interface for basic tasks.

The `RunTask` operation will support these maintenance commands. The CDDL framework will make no interpretation of the commands contained in the task. Each task will be a named task for tracking by external infrastructure, as well as provide the basic

content of the task. Optionally, the task can contain structured content from any derived namespace.

```
<cmp:taskActionRequest>  
  <cmp:taskName>xsd:string</cmp:taskName>  
  <cmp:taskContents>{any}</cmp:taskContents>  
</cmp:taskActionRequest>
```

All tasks will execute synchronously and in response to a task command, a response **MUST** be generated. A fault is **NOT** expected to be transmitted from this request, unless it is a fault of the communications or security infrastructure. The result of the task follows that of the input with the addition of an **OPTIONAL** integer status code.

```
<cmp:taskActionResponse>  
  <cmp:taskName>xsd:string</cmp:taskName>  
  <cmp:taskStatusCode>xsd:integer</cmp:taskStatusCode>  
  <cmp:taskOutput>{any}</cmp:taskOutput>  
</cmp:taskActionResponse>
```

4.5 Faults

All of the operations defined in the component model may return faults as an indication that an error has occurred. All faults are based on the WS-BaseFaults specification. Faults are raised in response to errors at resource endpoints or in the communication infrastructure or other portions of the distributed system. Faults can be transmitted in response to operations on the component endpoint or sent as part of notification events. Service implementations may also implicitly raise SOAPFault faults, as that is inherent in most implementations.

4.5.1 Fault Categories

There is a base fault `DeploymentFault`, from which all other faults are derived. This fault class directly overrides the `WS-BaseFault` fault type. It adds extra information useful for debugging and managing distributed systems. This items are reflective of the lessons learned in web services development [Loughran02], namely that extra information such as hostname and process is essential for locating which process among many has failed on a clustered system.

All interfaces must declare that they raise these `DeploymentFault` instances, rather than list the specific faults. This is to provide forward extensibility. If an implementation has a fault state whose meaning matches that of the predefined fault, the predefined fault **MUST** be thrown. If this predefined fault has standard elements for embedded fault information, the implementation **SHOULD** fill them in. The implementation **MAY** add implementation-specific data within the `ExtraData` element of the fault, to supplement this information. This extra data **MUST** not add new types to the XML namespaces of this deployment data. The XML schema and semantics of this extra data should be documented.

If an existing fault type is not suitable, implementations **MAY** create new fault types. If an implementation creates new fault types, these **MUST** extend the existing fault types

which operations are declared as throwing, which effectively means that they must extend `DeploymentFault`. These new faults **MUST NOT** change the XML schemas of CDDLM, and they **MUST** be in a new namespace. The new faults and XML content **SHOULD** be publicly documented.

If an implementation adds new operations or properties at the existing endpoints, these new operations **MAY** raise whatever faults they see fit, within the constraints of the `WS-BaseFault` specification. Again, the implementation must not add new types to the CDDLM namespaces.

4.5.1.1 *Transport faults*

Transport faults will inevitably be raised as the appropriate fault for the system. For example, the Apache Axis SOAP client raises `AxisFault` faults for all SOAP events, wrapping stack trace and even HTTP Fault data within the fault as DOM elements. Microsoft .NET WSE has a similar fault class.

4.5.1.2 *Relayed Faults*

Relayed faults are those received by the far end and passed on. They may be `WS-BaseFault` Faults; HTTP error codes, SOAP faults, native language faults wrapped as `SOAPFault`s, or predefined deployment faults. `WS-BaseFault` uses fault nesting for relaying faults; however, all faults must be a derivative of `WS-BaseFault`. This is addressed by defining a new `WS-BaseFault` derivative, a `WrappedSOAPFault`. This type is actually an extension of `DeploymentFault`. This fault can nest any received `SOAPFault`, with an element containing the received XML data. Well-known elements in this fault data (such as the Apache Axis stack trace and HTTP fault code) should be copied into any fields in the main fault that fill the same role.

4.5.2 Fault Security

Sites offering deployment services, may, for security reasons, wish to strip out some information, such as stack trace data. Implementations should provide a means to enable such an action prior to transmitting faults to callers.

Host name and process information may be viewed as sensitive, yet again, this is exceedingly useful to operations. Implementations may provide a means to disguise this information, so that it does not describe the real hostname or process ID of a process, but instead pseudonyms that can still be used in communications with any operations team.

4.5.3 Internationalization

The `WS-BaseFault` specification makes no statement upon which language error descriptions are described. If an implementation can return descriptions in one language, it must use `xml:lang` attributes to indicate the language of a description. Multiple descriptions, in different languages may be included. The client application should extract the description(s) whose language is the nearest match to that of the client.

4.5.4 DeploymentFault

This type represents any fault thrown by a component or the deployment infrastructure. All endpoint operations **MUST** declare that they throw this fault, and must not explicitly declare any derivative faults that they may throw.

Element	Type	Description
Host	xsd:string	Hostname or pseudonym.
Process	xsd:string	Any process identifier suitable for diagnostics.
ExtraData	unboundedXMLAnyNamespace	Extra fault data.
Component	xsd:string	Path to component raising the fault.
Stack	stringListType	Optional stack trace.

Implementations SHOULD include a component reference if it is known. Implementations SHOULD also include hostname and process information. Process information may be a low-level identifier (such as an operating system process ID), or it may be some application specific identifier. Its role is merely to distinguish which process amongst many in a load-balanced implementation raised the fault.

4.5.5 Component Faults

Most faults thrown by components are of the type `ComponentFault` which does no extension of the base `DeploymentFault`. The one exception is the fault thrown due to errors in state transition operations.

When a lifecycle state transition action is invoked, the system may fail to properly achieve the next state. As defined by that action, it will generate a `StateActionFault` to indicate failure of the action. In addition to the failure, the fault will indicate the extended state information that would normally be part of a status query directly inline within the fault. The fault may also include other data as defined by the `DeploymentFault`. This fault is defined by the following XML representation.

```
<cmp:StateActionFault>
  <wsbf:Timestamp>xsd:dateTime</wsbf:Timestamp>
  <wsbf:OriginatorReference>
    wsa:EndpointReferenceType
  </wsbf:OriginatorReference> ?
  <wsbf:ErrorCode dialect="anyURI">xsd:string</wsbf:ErrorCode> ?
  <wsbf:Description>xsd:string</wsbf:Description> *
  <wsbf:FaultCause>wsrf-bf:BaseFault</wsbf:FaultCause> *
  <cmp:ExtendedState>{any}</cmp:ExtendedState>
</cmp:StateActionFault>
```

4.6 Events

During the course of a deployment, many different lifecycle events will occur. Components are able to notify and receive notifications for these events, as well as forward these events to the Deployment API or other registered listeners. The CDL document can further describe the deployment process by declaring deployment events and event handlers.

Any deployment component in the deployment graph is a candidate for generating and receiving events. If the root of the deployment graph is specified as the recipient, the system EPR will be notified.

4.6.1 Event Interfaces

The WS-BaseNotification specification defines interfaces for producers and consumers of events. Components may be both producers and consumers simultaneously. It is RECOMMENDED that all components implement both the NotificationProducer and NotificationConsumer interfaces of WS-BaseNotification. An implementer MAY choose not to support events, but should realize that the developed component will not be able to be generally used outside of the specific deployment scenario it was created for.

These WS-BaseNotification interfaces require the implementation of the Notify, Subscribe and GetCurrentMessage operations. Documentation on these operations can be found in [WS-BF] and its accompanying WSDL.

4.6.2 Event Format

The implementation of the following events MUST be done using the WSDM Event Format, which extends the WS-BaseNotification specification. Additional notification mechanisms MAY also be used without restriction. The component model requires that a component MUST create a notification subscription for each event defined. As described in the Deployment API, there can be no guarantee of fault tolerant event subscriptions. Implementations of the Component Model MAY choose to include WS-Policy or other metadata to inform components how to renew subscriptions in the event of failures.

All events used by the Component Model are represented by a WSDM ManagementEvent [MUWS]. Each event is uniquely defined by an event identifier that defines the type of the management event and the resource which is the source of the event. The Component Model event information is added as sub-elements of the WSDM event.

4.6.2.1 Lifecycle Events

The lifecycle event used by the Component Model is an implementation of the WSDM State Capability event. Within the core MUWS event, the resource identifier and endpoint address are accompanied by an event identifier and a timestamp of the event. The event identifier is a URI which uniquely defines the type of event being transmitted, defined as <http://www.gridforum.org/cddlm/components/2005/02/events/Lifecycle>.

The Component Model defines the event message to contain a `cmp:LifecycleTransition` element as its “message content” defined by MUWS. This provides an indication of the lifecycle transition which has occurred to the consumer of the event. The message content MUST also contain a `muws-p2-xs:Situation` element. The only situation category supported by this specification is the `cmp:LifecycleSituation` declaring a change in CDDLM lifecycle state has occurred.

An example event would appear as in the following XML fragment.

```

<muws-p1-xs:ManagementEvent>
  <muws-p1-xs:ReportTime>2005-03-01T01:56:00Z</muws-p1-xs:ReportTime>
  <muws-p1-xs:EventId>http://www.gridforum.org/cddlm/components/2005/02/events/Lifecycle
</muws-p1-xs:EventId>
  <muws-p1-xs:SourceComponent>
    <muws-p1-xs:ResourceId>xsd:anyURI</muws-p1-xs:ResourceId>
  </muws-p1-xs:SourceComponent>
  <muws-p2-xs:Situation>
    <muws-p2-xs:SituationCategory>
      <cmp:LifecycleSituation/>
    </muws-p2-xs:SituationCategory>
  </muws-p2-xs:Situation>
  <cmp:LifecycleTransition>
    <muws-p2-xs:StateTransition Time="2005-03-01T01:54:30Z">
      <muws-p2-xs:EnteredState>
        <cmp:RunningState/>
      </muws-p2-xs:EnteredState>
      <muws-p2-xs:PreviousState>
        <cmp:InitializedState/>
      </muws-p2-xs:PreviousState>
    </muws-p2-xs:StateTransition>
  </cmp:LifecycleTransition>
</muws-p1-xs:ManagementEvent>

```

Events can be derived from this event type as defined by the WSDM MUWS specification, which allows both attributes and elements to be added to the main lifecycle event element. This is done through implementing an event specific instance of the muws-p2-xs:Situation element and extending its categorization.

In order to register for a Component Model event, a potential subscriber MUST use the WSDM StateCapability topic. The Component Model does not define any CDDLM specific topic for lifecycle events. These topics are defined as described in WS-Topics [WS-T] and WSDM MUWS [MUWS].

4.6.2.2 *Property Change Events*

When a property of a deployment component changes, it can notify any listeners of the change. There is no Component Model specific event defined. The property value change specification provided by WS-ResourceProperties [WS-RP] describes a detailed notification event to be used in this case to support transmission of the old and new element values.

4.6.3 **Event Registration**

Within the CDL document declaring components, event notifications will also be declared. All event declarations will take the same form. A deployment component will contain a CDL property that declares which event to register a notification for and the target of the notification. The component model requires that a deployment component MUST support the automatic registration of these event handlers at the time of component instantiation.

If it is desired to notify more than one component of the occurrence of a specific event, the property MAY be multiply defined with different target components. All target components MUST be deployment components or component delegates. They MUST NOT be language components or deployment components that are children of component delegates.

In the event of component failures, any restarted components SHOULD ensure that prior event registrations are restored.

4.6.3.1 *OnInitialized*

When a component transitions to the initialized state, it can notify one or more components that this has occurred. The target attribute will be declared using the CDL reference attribute. This reference is an XPath expression as defined in the CDL Language specification. This expression can be used to find the endpoint reference of the target component, using reference resolution as defined in Section 6.

In addition to notification, the component can also execute a state operation on the remote component. This state operation is indicated by use of one of the process attribute properties other than the “notify” attribute value. Additionally, this notification can be cancelled by indicating a value of the property to be false.

```
<cmp:OnInitialized process="initialize|run|terminate|destroy|notify"
  target="cdl:ref">
  [true|false|null]
</cmp:OnInitialized>
```

4.6.3.2 *OnRunning*

When a component has begun the running state, it can notify one or more components that this has occurred.

```
<cmp:OnRunning process="initialize|run|terminate|destroy|notify"
  target="cdl:ref">
  [true|false|null]
</cmp:OnRunning>
```

4.6.3.3 *OnFailed*

When a component has failed, it can notify one or more components that this has occurred.

```
<cmp:OnFailed process=" initialize|run|terminate|destroy|notify"
  target="cdl:ref">
  [true|false|null]
</cmp:OnFailed>
```

4.6.3.4 *OnTerminated*

When a component has terminated, it can notify one or more components that this has occurred.

```
<cmp:OnTerminated process="" initialize|run|terminate|destroy|notify"
  target="cdl:ref">
[true|false|null]
</cmp:OnTerminated>
```

4.6.4 Property Change Notifications

During the lifecycle of a component, it may be necessary to respond to other changes in the deployment, outside of lifecycle events. The WSRF structure allows consumers to register for notifications of changes in the properties of a stateful resource. We will enable this notification to be setup in the CDL document using the following syntax.

```
<cmp:OnChange property="xsd:QName" ?
  notify="cdl:ref">
```

This will enable a component to notify an element at the cdl:ref destination of the specific property change. If the property attribute is not used, then the component will register for notifications for all properties defined for the component.

4.6.5 Fault Notifications

A component may also wish to be notified that a fault was thrown by another component or deployment operation. It may subscribe for this notification by using the following syntax in the CDL Language

```
<cmp:OnFault notify="cdl:ref"
  faultName="xsd:QName" ?
  faultType="xsd:QName" ? >
```

The fault handling can optionally be scoped by including the fault name or the fault type in the handler's declaration

4.7 Control Flow

Many deployments follow simple rules, such as deploy these components in order. Or, deploy these components in any order you see fit. In order use a declarative means of creating our deployment description and retaining a simple syntax, we will add a few basic process elements. Basic flow control will allow an implementation to make simple decisions based on the state of deployment components or attributes. We will choose to use a very small subset of standard control primitives to aid in this task

The following sections describe the definition and syntax of these introduced constructs. The syntax borrows heavily from WS-BPEL and other languages for process control. It is not a complete calculus, nor does it need to be. It is minimally defined with only the operations necessary for simple control of deployment. For a discussion of complex deployment scenarios, plus consult Section 5.

When a system is defined, all top level objects are implicit children of a root system node in the graph, which can be equated to the system EPR. When a flow control element is

applied, it will be done as a child of a property node or the `cdl:system` declaration which will make this a child of the root system object. When applied, a flow control element will apply to all children of the node applied to.

As an example, if we wish to have several components deployed in order, we can use the `<cmp:sequence>` element defined below. As this node is declared at the root of the `cdl:system` structure, it will encapsulate the other declared elements within its control body.

```
<cdl:cdl
  targetNamespace="http://example.org/webapp-template">
<cdl:system>
<cmp:sequence lifecycle="initialization"/>
<WebApplication>
  <cmp:CodeBase>http://server/file.jar</cmp:CodeBase>
  <cmp:CommandPath>ApacheDeployer</cmp:CommandPath>
  <port>80</port>
  <hostname>www.example.org</hostname>
</WebApplication>
<Database>
...
</Database>
<TransactionMonitor>
...
</TransactionMonitor>
</cdl:system>
</cdl:cdl>
```

All flow control elements depend on lexical order and hierarchy of the declared `cdl:system` to apply their control functions. A `<cmp:sequence>` element as used above uses the lexical order of the declaration to order the deployment. By default, without the presence of a flow control element, an implementer of CDDL is free to apply any policy to ordering component operations. Once a flow control declaration has been made, it applies to all nodes in the graph that are peers or children of the element in which it is declared, effectively setting a scope of policy. That policy can be changed on more restrictive scopes, such as through redefinition on one of a node's children, but cannot be redeclared as a child of the same element. An implementer of CDL SHOULD ignore any redeclaration. For a comprehensive discussion of deployment scope, please consult Section 5.1.

4.7.1 Sequence

The sequence operator defines a structure for a deployment operation to be done in sequential order. As elements are defined in the CDL document, their lexical order will be used to define this deployment order. A sequential operation MUST be done as a synchronous, blocking operation. As one element is being initialized, the system MUST wait for initialization to complete before initialization of the next element.

The sequence operator also allows a value to be defined. This value is either a boolean value indicating whether to apply the operator or not, or a null value equivalent to a true definition.

```
<cmp:sequence lifecycle="initialization|execution|termination">
  [true|false|null]
</cmp:sequence>
```

4.7.2 Reverse

The reverse operator defines a structure for a deployment operation to be done in the reverse of sequential order. This element will synchronously execute the deployment operation in the reverse of the declared lexical order.

```
<cmp:reverse lifecycle="initialization|execution|termination">
  [true|false|null]
</cmp:reverse>
```

4.7.3 Flow

It is often not important what order operations are taken. The flow operator allows the deployment process to be executed concurrently, irrespective of order of declaration. Once all of the elements declared within the scope of the flow have completed their operation, that segment of the operation is considered complete. Thus, the flow operator does not imply that this operation can be done at any time during the deployment, only that its elements do not require ordered, synchronized operation.

```
<cmp:flow lifecycle="initialization|execution|termination">
  [true|false|null]
</cmp:flow>
```

4.7.4 Wait

Many times it is important to have some portion of a deployment delay while another continues. The wait operator allows a portion of the deployment process to be paused for a fixed period of time, or until some absolute time. Within the scope of a sequence, it will delay all pending elements of the sequence. Within a flow, it will pause the execution of all elements lexically declared after the wait element.

```
<cmp:wait lifecycle="initialization|execution|termination"
  ( duration="xsd:duration" | until="xsd:dateTime" ) />
```

4.7.5 Switch

The switch operator is designed to be used to control branches of execution. The switch operator allows a choice to be made based on some property of the deployment graph. The condition element must be a valid XPath calculation on declared properties returning a boolean value. The switch operation transfers execution of the deployment process to the node referred in the matching case. It does not allow one component to be executed instead of another. It is ONLY for changing the flow of execution. All component objects declared in a CDL document MUST be operational.

```
<cmp:switch lifecycle="initialization|execution|termination">
  <cmp:case condition="bool-expr"> +
    <cdl:ref ref="cdl:ref"/>
  </cmp:case>
  <cmp:otherwise>
```

```
<cdl:ref ref="cdl:ref"/>
  </cmp:otherwise>
</cmp:switch>
```

4.8 Component-specific message sets and messages

Any other message can also be processed by the same endpoint. When the extra messages are to be used between component instances, it is prudent to probe for the class of the destination endpoint before sending messages, particularly a long set of messages. It is RECOMMENDED that a component implement the WSDM ManageabilityCharacteristics capability in order to provide this type of discovery. This WSDM capability defines the ManagementCapability resource property.

Note that the notion of *change* is different in an XML based infrastructure than classic compile-time binding. Changes to existing messages may not be significant from a backwards compatibility perspective, if the recipient can still provide the semantics that a caller expects. However, a caller built against a recent version of a component may encounter problems when calling an older version. For this reason, it is critical to add new URIs for message sets whenever the WSDL or supported payload of a message changes in a way that could break systems in such a way.

Components SHOULD NOT change their set of supported interfaces during their life, as this prevents users of a component relying on the results of this call for any period of time.

5 Reference Resolution

Many properties in the CDL Language use the `cdl:ref` attribute to define their target location. Once a component is deployed, this reference must be translated into an actual resource identifier. A component reference will translate into a component EPR. A property reference will translate into its component owner's EPR and the attribute resource name under that EPR.

If a component property is set to a CDL reference, that property should be flagged by the component. Before the property is used, the component MUST retrieve the value of the property. Currently, the only available means to perform this translation is for a component to execute the `resolve()` function of the system EPR as defined in the Deployment API. Once a component retrieves the EPR of the destination, it can make a request to that EPR for the resource property of the target endpoint.

Once the property value is retrieved, it MUST be substituted for the reference. The component also MUST block whatever step it is about to perform and wait for the reference resolution process. It is RECOMMENDED that a component scan its property list before taking any deployment action that would require the use of the property. If it is unknown when the property is used, the component SHOULD retrieve this value when the component is initialized.

This process requires that the system EPR is able to handle the scale of potential resolution requests. Also, each component must have some way to discover the system EPR. We have not defined these mechanisms here.

6 Using the Component Model

Each component object within a deployment will transition through several states during its operation. The CDDL framework must be able to coordinate those transitions as well as provide an assessment of the system's state through aggregation of the states of the individual components that comprise the system. This section of the document provides comprehensive semantics for defining the state of a deployment, rules for using the component model, and examples for use.

6.1 Scope

The properties contained in a CDL configuration document typically imply very little about the form of a deployment, rather they describe the elements to deploy. As described earlier, by default all components are top level peers of each other. Some of the syntax elements described in the prior chapter can be used to change this behavior and create deployment dependencies which will force order to the deployment.

6.1.1 Delegates

If a component is a delegate, then the deployment structure begins to form as a delegate implies a parent-child relationship among subcomponents. In the diagram below, the independent components, Components A and B, and the component delegate, Component C, exist at the top level of the graph. Components D and E in the diagram are children of Component C.

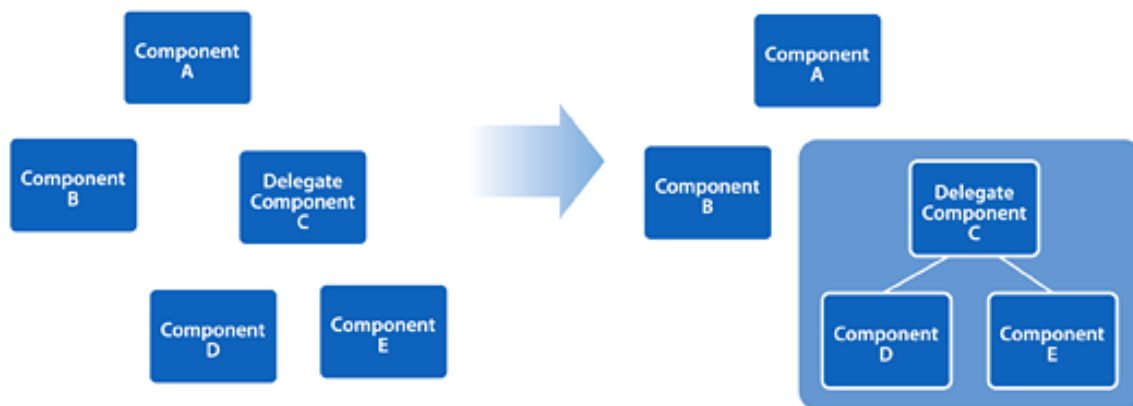


Figure 7. Deployment Scope.

When a component is declared as a delegate, its children are removed from scope of other components not declared as children of the delegate. This restricts further dependencies that are created within the graph to only exist between components in the same relative

scopes. If Component B in the diagram had a dependency on Component E, it MUST declare the dependency on Component C and trust that the delegate would enforce the desired behavior. In this sense, Component C becomes responsible for the aggregate lifecycles and dependencies of the objects it is a delegate for.

These dependencies includes lifecycle events and notifications, but not properties or property change notifications. If Component B needed a property value from Component D, it would not have to ask Component C for that value. But, it is an error for Component B to register for a lifecycle event on Component D. Further, if a flow control element were applied to this scenario, such as a sequence, Components A, B, and C would activate sequentially, but this flow element has no bearing on the process used to activate Components D and E. Only Component C has authority to activate those components.

There is no restriction on the depth of delegation allowed in this model. If a component is declared to be a child of a component delegate, it may also be a delegate itself, further responsible for its own children. Each level of delegation provides a more restrictive scope in which to apply the rules of deployment. Implicitly, if Component E in the above diagram were to have its own children F and G, Component D would not be allowed to imply any dependencies upon F or G, only through D itself.

The lifecycle of component delegates and their children are shared such that they appear to be identical to other components and the system. Whenever a delegate is created, its children MUST be created at the same time, in the order of definition within the CDL document. The delegate will not complete the creation action until all of its children have done so. This same rule applies to initialization and execution of the delegate. However, when the delegate fails or is terminated, its children MAY be acted upon asynchronously, allowing the delegate to notify failure or termination right away. The delegate SHOULD, though, propagate the failure or termination as with creation, in the lexical order of definition.

6.1.2 Document Structure

When the CDL document is created, the property lists may declare components all at one level as children of the `cdl:system` element, or may have some other lexical hierarchy. In the first case, the deployment graph would look like a simple list '(A B C D E), where the second would appear to be a tree or nested list '(A '(B C) D) E). In the strict syntax of CDL, property lists are trees and would require intermediate nodes to represent as this nested list, this syntax is used solely to show scope in a condensed fashion.

When flow control elements are added to the document, this hierarchy becomes important. In the first example, if a sequence element were declared as a child of the system element, the deployment would proceed as A->B->C->D->E. If the sequence element were declared as a child of one of the other elements, the deployment would happen at random, as the default behavior of the system is to deploy children of the system element as if a flow element were present.

This is due to the scope of application of the flow control element. When a flow control element is used, it applies to the children of the element it is declared upon. In the second example, if a sequence element were declared as a child of element A, then both the sequences A->B->C->D->E and E->A->B->C->D would be valid.

By applying flow control elements further into the tree, successively more restrictive scopes can be created. It is possible to redeclare a flow control element on a child element of a node where one was already declared. In the second example, declaring a flow element as a child of element A instead of a sequence element will not change the random order policy effected by the cdl:system node, but will instead mark a scope of flow control that may be used by other elements. This declaration would also ensure that if this element were imported into another CDL configuration document, it would enforce its desired behavior, rather than be at the whim of the declaration in the new document.

6.1.3 Flow Control Example

To understand the application of flow control and scope, we will look at an example deployment descriptor. In this example, a system is being deployed which contains an instance of the MySQL database, Apache Tomcat, and a web application JPetstore that requires a language configuration script to be run within the database in order to configure the application. In this example, the <cmp:sequence> element ensures that the four components are deployed and initialized in the declared order.

```
<cdl:system>
  <cmp:sequence lifecycle="initialization"/>
  <mysql/>
  <confLanguage/>
  <tomcat/>
  <jpetstore/>
</cdl:system>
```

When the deployment is terminated, there is no order specified for this portion of the lifecycle. If it is desired to terminate in the reverse order of the deployment, we can simply add a <cmp:reverse> element to the system definition. But, if we want to destroy confLanguage before mysql and jpetstore before tomcat, more precise control is needed over the order of the termination sequence. Instead, a deployment event can be used to enforce the desired order. An event such as <cmp:OnTerminated> allows definition of the exact sequences or simply notify other components of events.

```
<cdl:system>
  <cmp:sequence lifecycle="initialization"/>
  <mysql/>
  <confLanguage>
    <cmp:OnTerminated process="terminate" target="/mysql"/>
  </confLanguage>
  <tomcat/>
  <jpetstore>
    <cmp:OnTerminated process="terminate" target="/tomcat"/>
  </jpetstore>
</cdl:system>
```

This CDL indicates that the initialization should occur sequentially, but the termination has dependencies. confLanguage and jpetstore can be terminated in whatever order, but

the system should wait for confLanguage to notify the mysql component when it has successfully terminated, and wait for jpetstore to notify tomcat. If either jpetstore or confLanguage fail to terminate, the CDDL system should take action to resolve the failure.

6.2 Synchronization

Most operations of a deployment are done asynchronously. However, several of the properties contained in a CDL configuration document imply the need for synchronicity or wait states. Due to this, the basic hierarchy defined in the document can be modified further to reflect the order required by synchronization.

6.2.1 Events

Using events, one can declare that once Component A is running, it should instruct Component B to begin running. This can allow a deployment to be ordered and walked through its transitions in a simple and declarative manner. One could also instruct Component B that in the event it fails, it should notify Component A of this state transition. Then Component A could invoke some behavior such as a restart of Component B, or termination of itself and/or Component B.

Using this approach, we can create a deployment graph that is not singly linked, such as in a parent-child relationship like delegation. Each deployment action and event can independently link one component to one or more dependent components. If we were to expand the diagram from above to include this event example, it would appear as follows. This diagram is simple, but shows that the evolving graph has a chance to exhibit cyclic behaviors.




Figure 8. Deployment Event Structure.

The notification event is defined to be performed asynchronously, with no expectation from Component B that Component A will take any specific action. However, it is

known, a priori, that once Component A is running, it will inform Component B to run. An implementation of the component model **MUST** enforce this behavior. This requires that Component B waits for Component A to enter the running state and that no other deployment process tells Component B to execute.

As an example, if Components A, B and C were to be declared as children of a flow element, they would by default initialize concurrently. If Component B were declared to initialize Component C after it has initialized, then Components A and B would be initialized concurrently, but Component C would only be initialized after B has initiated the action.

Applying a sequence control flow element is functionally equivalent to declaring an event handler chain among the processes involved in the sequence. It is not required in an implementation of the component model to realize this structure in that fashion. However, one must recognize that if a sequence operation is declared and a CDL declaration inserts a notification within the sequence, the sequence will still continue in its own manner. If the consumer of the event causes other parts of the graph to change state, the implementer must recognize that state change events may not fire on those portions of the graph at the expected interval. Thus, the implementation of the sequence should be made resilient to this case.

6.2.2 Synchronization Points

It is often not important what order operations are taken. The flow operator allows the deployment process to be executed independent of its lexical declaration. A flow operator, as all other flow control elements, implies a synchronization point at the end of its scope. If a node has children that are declared with the flow operator, the node will still synchronize to its peers if necessary even if its children are executing concurrently. Thus, the completion of a deployment command is a guarantee of synchronization.

As an example, a CDL document may declare a deployment as follows (note that most of the declaration has been omitted for simplicity):

```
<cdl:cdl
  targetNamespace="http://example.org/sync-template">
<cdl:system>
<cmp:sequence lifecycle="initialization"/>
<ComponentA/>
<ComponentB>
  <cmp:flow lifecycle="initialization"/>
  <ComponentD/>
  <ComponentE/>
</ComponentB>
<ComponentC/>
</cdl:system>
</cdl:cdl>
```

In this case, the sequence operator defines that all children of the `cdl:system` node should be executed in order. This operator is overridden by the flow operator applied to Component B. This creates a scope encompassing Components' B, D and E. It implies

that Component A must initialize before B before C. Components B, D and E can be initialized in any order, but only after Component A has completed its initialization. Finally, Component C can only initialize after Components B, D and E.

If instead, the subscope of Component B were modified with a sequence operator, Component B would not be considered initialized until after Component E had completed its initialization. This defines a synchronization point at the scope of the flow control operator rather than at the component itself. Similarly, the switch and wait operators create synchronization points at the node which they are declared.

6.2.3 Events and Synchronization Points

It is important to recognize that deployment events and the defined control flow primitives together could potentially cause indeterminate or “spaghetti” control directives. Control flow primitives as discussed above define a scope with an implicit synchronization point. If a deployment event is declared between elements that cross these scopes, it MUST NOT alter the order implied by the target scope. Within a scope, a deployment event defines or redefines ordering as described earlier.

As an example, a CDL document may declare a deployment as follows:

```
<cdl:cdl
  targetNamespace="http://example.org/event-template">
<cdl:system>
<cmp:flow lifecycle="initialization"/>
<ComponentA/>
<ComponentB>
  <cmp:sequence lifecycle="initialization"/>
  <ComponentD/>
  <ComponentE/>
</ComponentB>
<ComponentC>
  <OnInitialized process="initialize" target="/ComponentB/ComponentE"/>
</ComponentC>
</cdl:system>
</cdl:cdl>
```

In this case, the basic deployment is unordered. However, when Component C has initialized, it must initialize Component E. Component B’s declaration specifies that Component D and Component E must be initialized in order. Thus, to ensure the proper ordering, and not violate scope, Component C must notify Component B to initialize, which will then initialize Components D and E. This creates an ordering to the root flow operator such that Component A and Component C can be initialized in any order, but Component B must be initialized after Component C.

6.3 Aggregate System State

Though the aggregate state of the deployment components is reflected in the system state, each component within the system is NOT REQUIRED to exist within the same state. Deployments may be constantly in flux, with components failing during the operations of the system. The system MAY respond to the failure(s) without affecting the state of the

whole, but the following recommendations should be followed. If all components are in a particular state, the system **MUST** report its state similarly. If the system is transitioning to a state, such as from instantiated to initialized, the system **SHOULD** report its state as the prior state, instantiated, until all components have reached the new state, initialized.

The system **MUST** also invoke actions in a coordinated fashion for all components. If two components are defined in a system then both must be initialized before they are run and so on. If the objects fail to transition to the same states, then the framework **SHOULD** coordinate their termination. Thus, if an object fails to initialize, the system **SHOULD** report its state as failed as well. The CDDL framework **MAY**, however, choose to reset and resume deployment for the failed portion of the task, or fail the entire deployment phase. It is important to note that the specific state of the system can only be determined by querying the system EPR. Any use of component states by entities outside of the CDDL framework or the system EPR is undefined.

If deployment event notifications are declared such as OnFailed events, the framework **MAY** wait to observe the results of the event propagation. If the events are declared in the CDL document as process steps rather than as events, such as `<OnFailed process="terminate">`, the framework **SHOULD** wait for the actions to propagate.

CDDL also enables component delegates to aggregate the lifecycle of other components. Using this aggregation, the main component represents a single instance of lifecycle state to the system, the sub components being opaque to the system in this regard. A component delegate **MAY** manage transitions of its children in any way it sees fit. For example, a delegate could choose to automatically restart a failed child node, but not need to indicate the failure to the system. The delegate itself, though, **MUST** follow the semantics of a deployment component itself with respect to the rest of the system. In this way, a delegate can be used to encode complex rules or interdependencies between components but still enforce the basic semantics of the component model.

6.4 Controlling Deployments

With effective use of just events and delegation, many complex deployment scenarios could be accommodated. However, the goal of the component model is to enable loosely coupled, scalable patterns for deployment. Events are an important tool for deployment, but may lead to significant extra effort to manage and maintain. For example, if we wanted to deploy our example system by deploying some components A then B then C, we would need to declare two events. That is, one less than the number of components to deploy. If we had a deployment with hundreds of components, the event system would be unwieldy at best.

Using delegation, one could create a complete structure for a deployment. However, in a large deployment, this will tightly couple components' behavior. Further, much of the rules describing how to react to deployment events will be lost within the components themselves. By creating a chain of parents and children, there is a potential for the relationships to restrict the form of the deployment and cause unwanted side effects. It is

recommended to use delegation sparingly. There are many cases, though, in which the delegation model will help with deployment.

6.4.1 Event Handlers

In order to handle more complex deployment scenarios, components can be defined whose sole purpose is to handle events. This allows a component to take responsibility for handling an event and determining itself how to respond to state transitions. These components should either be defined outside of the scope of a deployment process flow, or have simple internal state transitions of its own. These globally scoped event handlers can then be used as utilities to handle more complex operations such as restart on failure, or immediate termination on failure.

As an example, a component may define a notification to be sent to another component in the event that it fails. This component will remain in the failed state until the event handler chooses its course of action. The event handler *MAY* choose to perform an internal action to reset or fix the resource, in which case the component *MUST* transition back to the state it was in prior to failing. Otherwise, the event handler should perform its actions and forward the event to the system.

7 Security Considerations

For security, the CDDLM Component Model relies on the security provided by the CDL specification for securing the description of a deployment. Web Services security mechanisms, including transport-level security and SOAP Message Security are presumed to be used in order to secure the connections between components and EPRs.

The components themselves are responsible for performing actions within a trusted host. It is important that proper trust is established between host nodes in the deployment. Operations may require elevated permissions in order to perform. The component nodes *MUST* ensure that those rights cannot be usurped or delegated. The maintenance action task listed is most susceptible to security issues, as it may be possible to use this command to run arbitrary programs on remote nodes.

8 Editor Information

Stuart Schaefer
Softricity, Inc.
27 Melcher Street
Boston, MA 02210
Email: sschaefer@softricity.com

9 Acknowledgements

The authors of this document would like to acknowledge the contributions, assistance and support of the following people and their respective companies: Dejan Milojicic, Steve Loughran, Jun Tatemura, Takashi Kojo, Peter Toft, Patrick Goldsack, Julio Guijarro, Vanish Talwar, Sandro Rafaeli, and Bryan Murray.

References

- [RFC2119] Bradner, S. *Key words for use in RFCs to Indicate Requirement Levels*. IETF RFC 2119, March 1997. <http://www.ietf.org/rfc/rfc2119.txt?number=2119>.
- [Loughran02] Loughran, S. *Making Web Services that Work*. HP Laboratories, TR-HPL-2002-274, 2002. <http://www.hpl.hp.com/techreports/2002/HPL-2002-274.html>.
- [Loughran04] Loughran, S. *WS-RF and CDDLM*. CDDLM Working Group document, 2004. https://forge.gridforum.org/projects/cddlm-wg/document/WS-RF_and_CDDLM/en/1.
- [CDDLM] Milojicic, D., et al. *Configuration, Description, Deployment and Lifecycle Management (CDDLM) Foundation*. http://forge.gridforum.org/projects/cddlm-wg/document/CDDLM_Foundation_Document/en/1.
- [XML-CDL] Tatemura, J. *XML Configuration Description Language Specification*. CDDLM Working Group Draft, 2004. <http://forge.gridforum.org/>.
- [CDDLM-API] Loughran, S. *Deployment API Draft*. CDDLM Working Group Draft, 2004. <http://forge.gridforum.org/>.
- [WS-A] Box, D., et al. *Web Services Addressing*. W3C Member Submission, August 2004. <http://www.w3.org/Submission/2004/SUBM-ws-addressing-20040810/>
- [WS-RP] Graham, S. et al. *Web Services Resource Properties 1.2*. OASIS Working Draft, March 2005. <http://docs.oasis-open.org/wsrf/2005/03/wsrf-WS-ResourceProperties-1.2-draft-06.pdf>
- [WS-N] Graham, S. et al. *Web Services Base Notification 1.2*. OASIS Working Draft, June 2004. <http://docs.oasis-open.org/wsn/2004/06/wsn-WS-BaseNotification-1.2-draft-03.pdf>.
- [WS-BF] Tuecke, S. et al. *Web Services Base Faults*. OASIS Working Draft, March 2005. <http://docs.oasis-open.org/wsrf/2005/03/wsrf-WS-BaseFaults-1.2-draft-04.pdf>.
- [WS-T] Vambenepe, W. *Web Services Topics 1.2*. OASIS Working Draft, July 2004. <http://docs.oasis-open.org/wsn/2004/06/wsn-WS-Topics-1.2-draft-01.pdf>.
- [MUWS] Vambenepe, W. *Web Services Distributed Management: Management Using Web Services (MUWS 1.0) Part 1*. OASIS Committee Draft, December 2004. <http://docs.oasis-open.org/wsdm/2004/12/cd-wsdm-muws-part1-1.0.pdf>.
Vambenepe, W. *Web Services Distributed Management: Management Using Web Services (MUWS 1.0) Part 2*. OASIS Committee Draft, December 2004. <http://docs.oasis-open.org/wsdm/2004/12/cd-wsdm-muws-part2-1.0.pdf>.

Appendix A – Component Object Definitions

A.1 XML Schema

```

<?xml version="1.0" encoding="utf-8"?>
<xsd:schema xmlns="http://www.gridforum.org/cddlm/components/2005/02"
xmlns:xsd="http://www.w3.org/2001/XMLSchema" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:s12="http://www.w3.org/2003/05/soap-envelope"
xmlns:cdl="http://www.gridforum.org/2004/12/CDDL/XML-CDL/1.0"
xmlns:wsp="http://schemas.xmlsoap.org/ws/2002/12/policy"
xmlns:wsa="http://schemas.xmlsoap.org/ws/2003/03/addressing" xmlns:wsc="http://docs.oasis-
open.org/wsr/2004/06/wsr-WS-BaseFaults-1.2-draft-01.xsd" xmlns:wsp="http://docs.oasis-
open.org/wsr/2004/06/wsr-WS-ResourceProperties-1.2-draft-01.xsd" xmlns:wsl="http://docs.oasis-
open.org/wsr/2004/06/wsr-WS-ResourceLifetime-1.2-draft-01.xsd" xmlns:wnt="http://docs.oasis-
open.org/wsn/2004/06/wsn-WS-BaseNotification-1.2-draft-01.xsd" xmlns:wstop="http://docs.oasis-
open.org/wsn/2004/06/wsn-WS-Topics-1.2-draft-01.xsd" xmlns:muws-p1-xs="http://docs.oasis-
open.org/wsdm/2004/12/muws/wsdm-muws-part1.xsd" xmlns:muws-p2-xs="http://docs.oasis-
open.org/wsdm/2004/12/muws/wsdm-muws-part2.xsd"
targetNamespace="http://www.gridforum.org/cddlm/components/2005/02" elementFormDefault="qualified"
id="components">
  <!-- ===== -->
  <xsd:annotation>
    <xsd:documentation>

      This is the XSD describing the types of the CDDL Component Model

      Version: 1.0

      The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL
      NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and
      "OPTIONAL" in this document are to be interpreted as described in
      RFC 2119.
      http://www.ietf.org/rfc/rfc2119.txt

    </xsd:documentation>
  </xsd:annotation>
  <!-- ===== -->
  <!-- BASE TYPES -->
  <!-- ===== -->
  <xsd:complexType name="unboundedXMLOtherNamespace">
    <xsd:sequence>
      <xsd:any namespace="##other" processContents="lax" minOccurs="0" maxOccurs="unbounded"/>
    </xsd:sequence>
  </xsd:complexType>
  <!-- low level type: XML with no validations, namespace = ##any -->
  <xsd:complexType name="unboundedXMLAnyNamespace">
    <xsd:sequence>
      <xsd:any namespace="##any" processContents="lax" minOccurs="0" maxOccurs="unbounded"/>
    </xsd:sequence>
  </xsd:complexType>
  <!-- low level type: XML with no validations, namespace = ##any -->
  <xsd:complexType name="unboundedXMLAnyNamespaceWhitespacePreserved">
    <xsd:sequence>
      <xsd:any namespace="##any" processContents="lax" minOccurs="0" maxOccurs="unbounded"
xml:space="preserve"/>
    </xsd:sequence>
  </xsd:complexType>
  <xsd:simpleType name="void">

```

```

<xsd:restriction base="xsd:string">
  <xsd:length value="0"/>
</xsd:restriction>
</xsd:simpleType>
<xsd:simpleType name="booleanOrNull">
  <xsd:union memberTypes="xsd:boolean void" />
</xsd:simpleType>
<!-- ===== -->
<!-- list of [string] -->
<!-- ===== -->
<xsd:complexType name="stringListType">
  <xsd:annotation>
    <xsd:documentation>
      This is a type representing a list of strings
    </xsd:documentation>
  </xsd:annotation>
  <xsd:sequence>
    <xsd:element name="item" type="xsd:string" maxOccurs="unbounded"/>
  </xsd:sequence>
</xsd:complexType>
<!-- ===== -->
<!-- STATE MODEL -->
<!-- ===== -->
<xsd:complexType name="lifecycleStateType">
  <xsd:complexContent>
    <xsd:extension base="muws-p2-xs:StateType"/>
  </xsd:complexContent>
</xsd:complexType>
<xsd:element name="UndefinedState">
  <xsd:annotation>
    <xsd:documentation>Undefined state.</xsd:documentation>
  </xsd:annotation>
  <xsd:complexType>
    <xsd:complexContent>
      <xsd:restriction base="lifecycleStateType"/>
    </xsd:complexContent>
  </xsd:complexType>
</xsd:element>
<xsd:element name="InstantiatedState">
  <xsd:annotation>
    <xsd:documentation>State representing the presence of a component instance.</xsd:documentation>
  </xsd:annotation>
  <xsd:complexType>
    <xsd:complexContent>
      <xsd:restriction base="lifecycleStateType"/>
    </xsd:complexContent>
  </xsd:complexType>
</xsd:element>
<xsd:element name="InitializedState">
  <xsd:annotation>
    <xsd:documentation>State in which a component has been properly initialized.</xsd:documentation>
  </xsd:annotation>
  <xsd:complexType>
    <xsd:complexContent>
      <xsd:restriction base="lifecycleStateType"/>
    </xsd:complexContent>
  </xsd:complexType>
</xsd:element>
<xsd:element name="RunningState">
  <xsd:annotation>
    <xsd:documentation>Operational state.</xsd:documentation>
  </xsd:annotation>
  <xsd:complexType>
    <xsd:complexContent>

```

```

    <xsd:restriction base="lifecycleStateType"/>
  </xsd:complexContent>
</xsd:complexType>
</xsd:element>
<xsd:element name="FailedState">
  <xsd:annotation>
    <xsd:documentation>State in which the component has failed either a lifecycle operation or its operation has failed.</xsd:documentation>
  </xsd:annotation>
  <xsd:complexType>
    <xsd:complexContent>
      <xsd:restriction base="lifecycleStateType"/>
    </xsd:complexContent>
  </xsd:complexType>
</xsd:element>
<xsd:element name="TerminatedState">
  <xsd:annotation>
    <xsd:documentation>State in which a component instance has been terminated.</xsd:documentation>
  </xsd:annotation>
  <xsd:complexType>
    <xsd:complexContent>
      <xsd:restriction base="lifecycleStateType"/>
    </xsd:complexContent>
  </xsd:complexType>
</xsd:element>
<xsd:element name="LifecycleTransition" type="muws-p2-xs:StateTransitionType">
  <xsd:annotation>
    <xsd:documentation>A representation of the state transition that has occurred.</xsd:documentation>
  </xsd:annotation>
</xsd:element>
<!-- ===== -->
<!-- PROPERTIES -->
<!-- ===== -->
<xsd:simpleType name="lifecycleProcessEnum">
  <xsd:restriction base="xsd:string">
    <xsd:enumeration value="initialization">
      <xsd:annotation>
        <xsd:documentation>Apply the process during the initialization phase of the deployment.</xsd:documentation>
      </xsd:annotation>
    </xsd:enumeration>
    <xsd:enumeration value="execution">
      <xsd:annotation>
        <xsd:documentation>Apply the process during the transition to the execution phase.</xsd:documentation>
      </xsd:annotation>
    </xsd:enumeration>
    <xsd:enumeration value="termination">
      <xsd:annotation>
        <xsd:documentation>Apply the process during the termination phase of the deployment.</xsd:documentation>
      </xsd:annotation>
    </xsd:enumeration>
  </xsd:restriction>
</xsd:simpleType>
<xsd:simpleType name="eventProcessEnum">
  <xsd:restriction base="xsd:string">
    <xsd:enumeration value="intialize">
      <xsd:annotation>
        <xsd:documentation>Execute the initialization operation.</xsd:documentation>
      </xsd:annotation>
    </xsd:enumeration>
    <xsd:enumeration value="run">
      <xsd:annotation>
        <xsd:documentation>Execute the run operation.</xsd:documentation>
      </xsd:annotation>
    </xsd:enumeration>
  </xsd:restriction>

```



```

    </xsd:annotation>
  </xsd:enumeration>
  <xsd:enumeration value="terminate">
    <xsd:annotation>
      <xsd:documentation>Execute the terminate operation.</xsd:documentation>
    </xsd:annotation>
  </xsd:enumeration>
  <xsd:enumeration value="destroy">
    <xsd:annotation>
      <xsd:documentation>Destroy the remote component.</xsd:documentation>
    </xsd:annotation>
  </xsd:enumeration>
  <xsd:enumeration value="notify">
    <xsd:annotation>
      <xsd:documentation>Notify the remote component.</xsd:documentation>
    </xsd:annotation>
  </xsd:enumeration>
</xsd:restriction>
</xsd:simpleType>
<xsd:complexType name="terminationRecordType">
  <xsd:annotation>
    <xsd:documentation>
      Termination Record
    </xsd:documentation>
  </xsd:annotation>
  <xsd:sequence>
    <xsd:element name="timestamp" type="xsd:dateTime"/>
    <xsd:element name="message" type="xsd:string" minOccurs="0"/>
    <xsd:element name="extraData">
      <xsd:complexType>
        <xsd:sequence>
          <xsd:any namespace="##other" processContents="lax" minOccurs="0" maxOccurs="unbounded"/>
        </xsd:sequence>
      </xsd:complexType>
    </xsd:element>
  </xsd:sequence>
</xsd:complexType>
<xsd:complexType name="commandPathType" mixed="true">
  <xsd:annotation>
    <xsd:documentation>A command path can contain a command and a list arguments.</xsd:documentation>
  </xsd:annotation>
  <xsd:sequence>
    <xsd:element name="path" type="xsd:string"/>
    <xsd:element name="args" type="xsd:anyType" minOccurs="0" maxOccurs="unbounded"/>
  </xsd:sequence>
</xsd:complexType>
<xsd:complexType name="componentStatusType">
  <xsd:annotation>
    <xsd:documentation>This is the status of a deployed component.</xsd:documentation>
  </xsd:annotation>
  <xsd:sequence>
    <xsd:element name="State" type="lifecycleStateType"/>
    <xsd:element ref="LifecycleTransition"/>
    <xsd:element name="ExtendedState" type="unboundedXMLAnyNamespace" minOccurs="0"/>
  </xsd:sequence>
</xsd:complexType>
<!-- ===== -->
<!-- COMPONENT ELEMENTS -->
<!-- ===== -->
<xsd:element name="ComponentReference" type="wsa:EndpointReferenceType">
  <xsd:annotation>
    <xsd:documentation>Every component will be accessible via its WS-RF EPR.</xsd:documentation>
  </xsd:annotation>
</xsd:element>

```

```

<xsd:element name="CodeBase" type="xsd:anyURI">
  <xsd:annotation>
    <xsd:documentation>Location of the component file assets.</xsd:documentation>
  </xsd:annotation>
</xsd:element>
<xsd:element name="CommandPath" type="commandPathType">
  <xsd:annotation>
    <xsd:documentation>The location or object within the assets of the component to use in activating the
component.</xsd:documentation>
  </xsd:annotation>
</xsd:element>
<xsd:element name="Delegate" type="booleanOrNull">
  <xsd:annotation>
    <xsd:documentation>A component which controls other components.</xsd:documentation>
  </xsd:annotation>
</xsd:element>
<xsd:element name="ComponentStatus" type="componentStatusType"/>
<xsd:element name="ComponentName" type="cdl:pathType"/>
<xsd:element name="ComponentDescription" type="unboundedXMLAnyNamespace"/>
<xsd:element name="ComponentProperties">
  <xsd:annotation>
    <xsd:documentation>The WSRF Resource Property document fragment of a component.</xsd:documentation>
  </xsd:annotation>
  <xsd:complexType>
    <xsd:sequence>
      <xsd:element ref="muws-p1-xs:ResourceId"/>
      <xsd:element ref="muws-p1-xs:ManageabilityCapability"/>
      <xsd:element ref="ComponentName"/>
      <xsd:element ref="ComponentDescription"/>
      <xsd:element ref="ComponentStatus"/>
      <xsd:element ref="wsnt:Topic" maxOccurs="unbounded"/>
      <xsd:element ref="wsnt:FixedTopicSet"/>
      <xsd:element ref="wsnt:TopicExpressionDialects" maxOccurs="unbounded"/>
    </xsd:sequence>
  </xsd:complexType>
</xsd:element>
<!-- ===== -->
<!-- DEPLOYMENT OPERATORS -->
<!-- ===== -->
<xsd:complexType name="flowControlType">
  <xsd:simpleContent>
    <xsd:extension base="booleanOrNull">
      <xsd:attribute name="lifecycle" type="lifecycleProcessEnum" use="required"/>
    </xsd:extension>
  </xsd:simpleContent>
</xsd:complexType>
<xsd:element name="sequence" type="flowControlType">
  <xsd:annotation>
    <xsd:documentation>A sequential deployment operator.</xsd:documentation>
  </xsd:annotation>
</xsd:element>
<xsd:element name="reverse" type="flowControlType">
  <xsd:annotation>
    <xsd:documentation>A reverse sequential deployment operator.</xsd:documentation>
  </xsd:annotation>
</xsd:element>
<xsd:element name="flow" type="flowControlType">
  <xsd:annotation>
    <xsd:documentation>A deployment operator for creating concurrent deployment
processes.</xsd:documentation>
  </xsd:annotation>
</xsd:element>
<xsd:element name="wait">
  <xsd:annotation>

```

```

<xsd:documentation>A deployment operator for blocking a deployment process for a period of
time.</xsd:documentation>
</xsd:annotation>
<xsd:complexType>
  <xsd:attribute name="duration" type="xsd:duration"/>
  <xsd:attribute name="until" type="xsd:dateTime"/>
</xsd:complexType>
</xsd:element>
<xsd:element name="switch">
  <xsd:annotation>
    <xsd:documentation>A deployment operator for branching deployment processes.</xsd:documentation>
  </xsd:annotation>
  <xsd:complexType>
    <xsd:sequence>
      <xsd:element ref="case" minOccurs="0" maxOccurs="unbounded"/>
      <xsd:element ref="otherwise"/>
    </xsd:sequence>
    <xsd:attribute name="lifecycle" type="lifecycleProcessEnum" use="required"/>
  </xsd:complexType>
</xsd:element>
<xsd:element name="case">
  <xsd:annotation>
    <xsd:documentation>A branch of a deployment process.</xsd:documentation>
  </xsd:annotation>
  <xsd:complexType>
    <xsd:all>
      <xsd:element ref="cdl:ref"/>
    </xsd:all>
    <xsd:attribute name="condition" type="xsd:string" use="required"/>
  </xsd:complexType>
</xsd:element>
<xsd:element name="otherwise">
  <xsd:annotation>
    <xsd:documentation>A default branch of a deployment process.</xsd:documentation>
  </xsd:annotation>
  <xsd:complexType>
    <xsd:all>
      <xsd:element ref="cdl:ref"/>
    </xsd:all>
  </xsd:complexType>
</xsd:element>
<!-- ===== -->
<!-- DEPLOYMENT EVENTS -->
<!-- ===== -->
<xsd:complexType name="eventRegistrationType">
  <xsd:simpleContent>
    <xsd:extension base="booleanOrNull">
      <xsd:attribute name="process" type="eventProcessEnum" use="required"/>
      <xsd:attribute name="target" type="cdl:pathType" use="required"/>
    </xsd:extension>
  </xsd:simpleContent>
</xsd:complexType>
<xsd:element name="OnInitialized" type="eventRegistrationType">
  <xsd:annotation>
    <xsd:documentation>Notify a component that the source has achieved the initialized
state.</xsd:documentation>
  </xsd:annotation>
</xsd:element>
<xsd:element name="OnRunning" type="eventRegistrationType">
  <xsd:annotation>
    <xsd:documentation>Notify a component that the source has achieved the running
state.</xsd:documentation>
  </xsd:annotation>
</xsd:element>

```

```

<xsd:element name="OnFailed" type="eventRegistrationType">
  <xsd:annotation>
    <xsd:documentation>Notify a component that the source has entered the failed state.</xsd:documentation>
  </xsd:annotation>
</xsd:element>
<xsd:element name="OnTerminated" type="eventRegistrationType">
  <xsd:annotation>
    <xsd:documentation>Notify a component that the source has entered the terminated
state.</xsd:documentation>
  </xsd:annotation>
</xsd:element>
<xsd:element name="OnChange">
  <xsd:annotation>
    <xsd:documentation>Notify a component that a property of the source has changed.</xsd:documentation>
  </xsd:annotation>
  <xsd:complexType>
    <xsd:attribute name="property" type="xsd:QName" use="optional"/>
    <xsd:attribute name="notify" type="cdl:pathType" use="required"/>
  </xsd:complexType>
</xsd:element>
<xsd:element name="OnFault">
  <xsd:annotation>
    <xsd:documentation>Subscribe to a component to be notified of the occurrence of a named
fault.</xsd:documentation>
  </xsd:annotation>
  <xsd:complexType>
    <xsd:attribute name="notify" type="cdl:pathType" use="required"/>
    <xsd:attribute name="faultName" type="xsd:QName" use="optional"/>
    <xsd:attribute name="faultType" type="xsd:QName" use="optional"/>
  </xsd:complexType>
</xsd:element>
<!-- ===== -->
<!-- OPERATION MESSAGES -->
<!-- ===== -->
<xsd:element name="initializeRequest" type="void">
  <xsd:annotation>
    <xsd:documentation>Request to move a component to the initialized state. Valid only if the component is in
the instantiated state.</xsd:documentation>
  </xsd:annotation>
</xsd:element>
<xsd:element name="initializeResponse" type="void"/>
<xsd:element name="runRequest" type="void">
  <xsd:annotation>
    <xsd:documentation>Request to move a component to the running state. Valid only if the component is in the
initialized state, or it is already running.</xsd:documentation>
  </xsd:annotation>
</xsd:element>
<xsd:element name="runResponse" type="void"/>
<xsd:element name="terminateRequest" type="void">
  <xsd:annotation>
    <xsd:documentation>Request to move a component to the terminated state. Valid only if the component is in a
state from which it can enter the terminate state, or it has already terminated.</xsd:documentation>
  </xsd:annotation>
</xsd:element>
<xsd:element name="terminateResponse" type="void"/>
<xsd:element name="taskActionRequest">
  <xsd:complexType>
    <xsd:sequence>
      <xsd:element name="taskName" type="xsd:string"/>
      <xsd:element name="taskContents" type="unboundedXMLAnyNamespace" minOccurs="0"/>
    </xsd:sequence>
  </xsd:complexType>
</xsd:element>
<xsd:element name="taskActionResponse">

```

```

<xsd:complexType>
  <xsd:sequence>
    <xsd:element name="taskName" type="xsd:string"/>
    <xsd:element name="taskStatusCode" type="xsd:integer" minOccurs="0"/>
    <xsd:element name="taskOutput" type="unboundedXMLAnyNamespace" minOccurs="0"/>
  </xsd:sequence>
</xsd:complexType>
</xsd:element>
<!-- ===== -->
<!-- FAULTS -->
<!-- ===== -->
<xsd:complexType name="DeploymentFaultType">
  <xsd:annotation>
    <xsd:documentation>Base fault for all CDDL M faults</xsd:documentation>
  </xsd:annotation>
  <xsd:complexContent>
    <xsd:extension base="wsbf:BaseFaultType">
      <xsd:sequence>
        <xsd:element name="Host" type="xsd:string" minOccurs="0"/>
        <xsd:element name="Process" type="xsd:string" minOccurs="0"/>
        <xsd:element name="ComponentName" type="cdl:pathType" minOccurs="0"/>
        <xsd:element name="Stack" type="stringListType" minOccurs="0"/>
        <xsd:element name="ExtraData" type="unboundedXMLAnyNamespace" minOccurs="0"/>
      </xsd:sequence>
    </xsd:extension>
  </xsd:complexContent>
</xsd:complexType>
<xsd:element name="ComponentFault" type="DeploymentFaultType"/>
<xsd:complexType name="StateActionFaultType">
  <xsd:annotation>
    <xsd:documentation>Fault to indicate a state transition action has failed.</xsd:documentation>
  </xsd:annotation>
  <xsd:complexContent>
    <xsd:extension base="DeploymentFaultType">
      <xsd:sequence>
        <xsd:element name="ExtendedState" type="unboundedXMLAnyNamespace" minOccurs="0"/>
      </xsd:sequence>
    </xsd:extension>
  </xsd:complexContent>
</xsd:complexType>
<xsd:element name="StateActionFault" type="StateActionFaultType"/>
<!-- Wrapped SOAP Fault -->
<xsd:complexType name="WrappedSOAPFaultType">
  <xsd:annotation>
    <xsd:documentation>SOAP Fault Fault</xsd:documentation>
  </xsd:annotation>
  <xsd:complexContent>
    <xsd:extension base="DeploymentFaultType">
      <xsd:sequence>
        <xsd:element name="SoapFault" type="s12:Fault"/>
      </xsd:sequence>
    </xsd:extension>
  </xsd:complexContent>
</xsd:complexType>
<xsd:element name="WrappedSOAPFault" type="WrappedSOAPFaultType"/>
</xsd:schema>

```

A.2 WSDL 1.1

```

<?xml version="1.0" encoding="UTF-8"?>
<wsdl:definitions xmlns:cmpw="http://www.gridforum.org/cddlm/components/2005/02/wsdl"
:wsdl="http://schemas.xmlsoap.org/wsdl/" :soap="http://schemas.xmlsoap.org/wsdl/soap/"
:soapenc="http://schemas.xmlsoap.org/soap/encoding/" :xsd="http://www.w3.org/2001/XMLSchema"

```

```

:wsa="http://schemas.xmlsoap.org/ws/2003/03/addressing" :wsbf="http://docs.oasis-
open.org/wsrp/2004/06/wsrp-WS-BaseFaults-1.2-draft-01.xsd" :wsrp="http://docs.oasis-
open.org/wsrp/2004/06/wsrp-WS-ResourceProperties-1.2-draft-01.xsd" xmlns:wsrpw="http://docs.oasis-
open.org/wsrp/2004/06/wsrp-WS-ResourceProperties-1.2-draft-01.wsdl":wsrl="http://docs.oasis-
open.org/wsrp/2004/06/wsrp-WS-ResourceLifetime-1.2-draft-01.xsd" xmlns:wsrlw="http://docs.oasis-
open.org/wsrp/2004/06/wsrp-WS-ResourceLifetime-1.2-draft-01.wsdl":wsnt="http://docs.oasis-
open.org/wsn/2004/06/wsn-WS-BaseNotification-1.2-draft-01.xsd" xmlns:wsntw="http://docs.oasis-
open.org/wsn/2004/06/wsn-WS-BaseNotification-1.2-draft-01.wsdl":muws-p1-xs="http://docs.oasis-
open.org/wsdm/2004/12/muws/wsdm-muws-part1.xsd"
targetNamespace="http://www.gridforum.org/cddlm/components/2005/02/wsdl">
<!-- ===== -->
<wSDL:documentation>

  This is the WSDL Describing the Component Model API for CDDL
  deployment components.

  The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL
  NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and
  "OPTIONAL" in this document are to be interpreted as described in
  RFC 2119.
  http://www.ietf.org/rfc/rfc2119.txt

</wSDL:documentation>
<!-- ===== -->
<!-- COMPONENT MESSAGES -->
<!-- ===== -->
<wSDL:message name="initializeRequest">
  <wSDL:part name="initializeRequest" element="cmp:initializeRequest"/>
</wSDL:message>
<wSDL:message name="initializeResponse">
  <wSDL:part name="initializeResponse" element="cmp:initializeResponse"/>
</wSDL:message>
<wSDL:message name="runRequest">
  <wSDL:part name="runRequest" element="cmp:runRequest"/>
</wSDL:message>
<wSDL:message name="runResponse">
  <wSDL:part name="runResponse" element="cmp:runResponse"/>
</wSDL:message>
<wSDL:message name="terminateRequest">
  <wSDL:part name="terminateRequest" element="cmp:terminateRequest"/>
</wSDL:message>
<wSDL:message name="terminateResponse">
  <wSDL:part name="terminateResponse" element="cmp:terminateResponse"/>
</wSDL:message>
<wSDL:message name="taskActionRequest">
  <wSDL:part name="taskActionRequest" element="cmp:taskActionRequest"/>
</wSDL:message>
<wSDL:message name="taskActionResponse">
  <wSDL:part name="taskActionRequest" element="cmp:taskActionRequest"/>
</wSDL:message>
<wSDL:message name="StateActionFault">
  <wSDL:part name="fault" element="cmp:StateActionFault"/>
</wSDL:message>
<!-- ===== -->
<!-- COMPONENT ENDPOINT -->
<!-- ===== -->
<wSDL:portType name="ComponentEndpoint" wsrp:ResourceProperties="cmp:ComponentProperties">
<!-- ===== WSRP Property Messages ===== -->
  <wSDL:operation name="GetResourceProperty">
    <wSDL:input name="GetResourcePropertyRequest" message="wsrpw:GetResourcePropertyRequest"/>
    <wSDL:output name="GetResourcePropertyResponse" message="wsrpw:GetResourcePropertyResponse"/>
    <wSDL:fault name="ResourceUnknownFault" message="wsrpw:ResourceUnknownFault"/>
    <wSDL:fault name="InvalidResourcePropertyQNameFault"
    message="wsrpw:InvalidResourcePropertyQNameFault"/>

```

```

</wsdl:operation>
<wsdl:operation name="GetMultipleResourceProperties">
  <wsdl:input name="GetMultipleResourcePropertiesRequest"
message="wsrpw:GetMultipleResourcePropertiesRequest"/>
  <wsdl:output name="GetMultipleResourcePropertiesResponse"
message="wsrpw:GetMultipleResourcePropertiesResponse"/>
  <wsdl:fault name="ResourceUnknownFault" message="wsrpw:ResourceUnknownFault"/>
  <wsdl:fault name="InvalidResourcePropertyQNameFault"
message="wsrpw:InvalidResourcePropertyQNameFault"/>
</wsdl:operation>
<wsdl:operation name="SetResourceProperties">
  <wsdl:input name="SetResourcePropertiesRequest" message="wsrpw:SetResourcePropertiesRequest"/>
  <wsdl:output name="SetResourcePropertiesResponse" message="wsrpw:SetResourcePropertiesResponse"/>
  <wsdl:fault name="ResourceUnknownFault" message="wsrpw:ResourceUnknownFault"/>
  <wsdl:fault name="InvalidSetResourcePropertiesRequestContentFault"
message="wsrpw:InvalidSetResourcePropertiesRequestContentFault"/>
  <wsdl:fault name="UnableToModifyResourcePropertyFault"
message="wsrpw:UnableToModifyResourcePropertyFault"/>
  <wsdl:fault name="InvalidResourcePropertyQNameFault"
message="wsrpw:InvalidResourcePropertyQNameFault"/>
  <wsdl:fault name="SetResourcePropertyRequestFailedFault"
message="wsrpw:SetResourcePropertyRequestFailedFault"/>
</wsdl:operation>
<wsdl:operation name="QueryResourceProperties">
  <wsdl:input name="QueryResourcePropertiesRequest"
message="wsrpw:QueryResourcePropertiesRequest"/>
  <wsdl:output name="QueryResourcePropertiesResponse"
message="wsrpw:QueryResourcePropertiesResponse"/>
  <wsdl:fault name="ResourceUnknownFault" message="wsrpw:ResourceUnknownFault"/>
  <wsdl:fault name="InvalidResourcePropertyQNameFault"
message="wsrpw:InvalidResourcePropertyQNameFault"/>
  <wsdl:fault name="UnknownQueryExpressionDialectFault"
message="wsrpw:UnknownQueryExpressionDialectFault"/>
  <wsdl:fault name="InvalidQueryExpressionFault" message="wsrpw:InvalidQueryExpressionFault"/>
  <wsdl:fault name="QueryEvaluationErrorFault" message="wsrpw:QueryEvaluationErrorFault"/>
</wsdl:operation>
<!-- ===== NotificationProducer Specific ===== -->
<wsdl:operation name="Subscribe">
  <wsdl:input name="SubscribeRequest" message="wsntw:SubscribeRequest"/>
  <wsdl:output name="SubscribeResponse" message="wsntw:SubscribeResponse"/>
  <wsdl:fault name="ResourceUnknownFault" message="wsntw:ResourceUnknownFault"/>
  <wsdl:fault name="SubscribeCreationFailedFault" message="wsntw:SubscribeCreationFailedFault"/>
  <wsdl:fault name="TopicPathDialectUnknownFault" message="wsntw:TopicPathDialectUnknownFault"/>
</wsdl:operation>
<wsdl:operation name="GetCurrentMessage">
  <wsdl:input name="GetCurrentMessageRequest" message="wsntw:GetCurrentMessageRequest"/>
  <wsdl:output name="GetCurrentMessageResponse" message="wsntw:GetCurrentMessageResponse"/>
  <wsdl:fault name="ResourceUnknownFault" message="wsntw:ResourceUnknownFault"/>
  <wsdl:fault name="InvalidTopicExpressionFault" message="wsntw:InvalidTopicExpressionFault"/>
  <wsdl:fault name="TopicNotSupportedFault" message="wsntw:TopicNotSupportedFault"/>
  <wsdl:fault name="NoCurrentMessageOnTopicFault" message="wsntw:NoCurrentMessageOnTopicFault"/>
</wsdl:operation>
<!-- ===== NotificationConsumer Specific ===== -->
<wsdl:operation name="Notify">
  <wsdl:input message="wsntw:Notify"/>
</wsdl:operation>
<!-- ===== Lifecycle Operations ===== -->
<!-- WS-RF Resource Lifetime: ImmediateResourceTermination -->
<wsdl:operation name="Destroy">
  <wsdl:input name="DestroyRequest" message="wsrlw:DestroyRequest"/>
  <wsdl:output name="DestroyResponse" message="wsrlw:DestroyResponse"/>
  <wsdl:fault name="ResourceUnknownFault" message="wsrlw:ResourceUnknownFault"/>
  <wsdl:fault name="ResourceNotDestroyedFault" message="wsrlw:ResourceNotDestroyedFault"/>
</wsdl:operation>

```

```
<!-- initialize a component -->
<wsdl:operation name="Initialize">
  <wsdl:input name="initializeRequest" message="cmpw:initializeRequest"/>
  <wsdl:output name="initializeResponse" message="cmpw:initializeResponse"/>
  <wsdl:fault name="StateActionFault" message="cmpw:StateActionFault"/>
</wsdl:operation>
<!-- run a component -->
<wsdl:operation name="Run">
  <wsdl:input name="runRequest" message="cmpw:runRequest"/>
  <wsdl:output name="runResponse" message="cmpw:runResponse"/>
  <wsdl:fault name="StateActionFault" message="cmpw:StateActionFault"/>
</wsdl:operation>
<!-- terminate a component -->
<wsdl:operation name="Terminate">
  <wsdl:input name="terminateRequest" message="cmpw:terminateRequest"/>
  <wsdl:output name="terminateResponse" message="cmpw:terminateResponse"/>
  <wsdl:fault name="StateActionFault" message="cmpw:StateActionFault"/>
</wsdl:operation>
<!-- maintenance operation -->
<wsdl:operation name="RunTask">
  <wsdl:input name="taskActionRequest" message="cmpw:taskActionRequest"/>
  <wsdl:output name="taskActionResponse" message="cmpw:taskActionResponse"/>
</wsdl:operation>
</wsdl:portType>
</wsdl:definitions>
```