



Cloud Application Management for Platforms Version 1.2

Committee Specification 01

15 May 2018

Specification URIs

This version:

<http://docs.oasis-open.org/camp/camp-spec/v1.2/cs01/camp-spec-v1.2-cs01.pdf> (Authoritative)
<http://docs.oasis-open.org/camp/camp-spec/v1.2/cs01/camp-spec-v1.2-cs01.html>
<http://docs.oasis-open.org/camp/camp-spec/v1.2/cs01/camp-spec-v1.2-cs01.doc>

Previous version:

<http://docs.oasis-open.org/camp/camp-spec/v1.2/csprd01/camp-spec-v1.2-csprd01.pdf>
(Authoritative)
<http://docs.oasis-open.org/camp/camp-spec/v1.2/csprd01/camp-spec-v1.2-csprd01.html>
<http://docs.oasis-open.org/camp/camp-spec/v1.2/csprd01/camp-spec-v1.2-csprd01.doc>

Latest version:

<http://docs.oasis-open.org/camp/camp-spec/v1.2/camp-spec-v1.2.pdf> (Authoritative)
<http://docs.oasis-open.org/camp/camp-spec/v1.2/camp-spec-v1.2.html>
<http://docs.oasis-open.org/camp/camp-spec/v1.2/camp-spec-v1.2.doc>

Technical Committee:

OASIS Cloud Application Management for Platforms (CAMP) TC

Chair:

Anish Karmarkar (anish.karmarkar@oracle.com), Oracle

Editors:

Jacques Durand (jdurand@us.fujitsu.com), Fujitsu Limited
Adrian Otto (adrian.otto@rackspace.com), Rackspace Hosting, Inc.
Gilbert Pilz (gilbert.pilz@oracle.com), Oracle
Tom Rutt (trutt@us.fujitsu.com), Fujitsu Limited

Related work:

This specification replaces or supersedes:

- *Cloud Application Management for Platforms Version 1.1*. Edited by Jacques Durand, Adrian Otto, Gilbert Pilz, and Tom Rutt. Latest version: <http://docs.oasis-open.org/camp/camp-spec/v1.1/camp-spec-v1.1.html>.

Test assertions for the normative statements in this specification are located at:

- *Cloud Application Management for Platforms (CAMP) Test Assertions Version 1.1*. Edited by Jacques Durand, Gilbert Pilz, Adrian Otto, and Tom Rutt. Latest version: <http://docs.oasis-open.org/camp/camp-ta/v1.1/camp-ta-v1.1.html>.

Abstract:

This document defines the artifacts and APIs that need to be offered by a Platform as a Service (PaaS) cloud to manage the building, running, administration, monitoring and patching of applications in the cloud. Its purpose is to enable interoperability among self-service interfaces to PaaS clouds by defining artifacts and formats that can be used with any conforming cloud and

enable independent vendors to create tools and services that interact with any conforming cloud using the defined interfaces. Cloud vendors can use these interfaces to develop new PaaS offerings that will interact with independently developed tools and components.

Status:

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

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

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

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

Citation format:

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

[CAMP-v1.2]

Cloud Application Management for Platforms Version 1.2. Edited by Jacques Durand, Adrian Otto, Gilbert Pilz, and Tom Rutt. 15 May 2018. OASIS Committee Specification 01. <http://docs.oasis-open.org/camp/camp-spec/v1.2/cs01/camp-spec-v1.2-cs01.html>. Latest version: <http://docs.oasis-open.org/camp/camp-spec/v1.2/camp-spec-v1.2.html>.

Notices

Copyright © OASIS Open 2018. All Rights Reserved.

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

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

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

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

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

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

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

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

Table of Contents

1	Introduction	10
1.0	IPR Policy	10
1.1	Overview	10
1.2	Purpose	10
1.3	Example (non-normative)	11
1.4	Non-Goals	12
1.5	Actors	12
1.6	Terminology	12
1.6.1	Term Definitions	12
1.6.1.1	CAMP Provider (Provider)	12
1.6.1.2	CAMP Consumer (Consumer)	12
1.6.2	Keywords, Conventions, and Normative Text	13
1.7	Notational Conventions	13
1.8	Specification Version	13
1.8.1	Backwards Compatibility	13
1.9	Normative References	13
1.10	Non-Normative References	14
2	Concepts and Types	16
2.1	Introduction	16
2.2	Resources	17
2.2.1	Platform	17
2.2.2	Assemblies	17
2.2.3	Components	18
2.2.4	Plans	18
2.2.5	Services	18
2.2.6	Collection and Factory Resources	18
2.2.7	Operations and Sensors	19
2.2.8	Resource Relationships	21
2.3	Deployment	21
2.4	Versions and Extensions	22
2.5	Parameters	24
2.6	CAMP Common Attribute Types	25
2.7	Representation Skew	26
2.8	External Management Resources	27
3	Application Management Lifecycle	28
3.1	Initial Platform Resources	28
3.2	Creating an Assembly from a PDP or Plan File	28
3.3	Creating an Assembly from a plan resource	29
3.4	Managing an Application Assembly	31
3.5	Removing Assemblies	31
4	Platform Deployment Package	32
4.1	PDP Package Structure	32
4.1.1	Supported Archive Formats	32

4.1.2 Validating Integrity	32
4.2 Plan Overview	32
4.2.1 Types	33
4.2.2 Requirement Specifications	33
4.2.3 Service Specifications	33
4.2.3.1 Shared Services	35
4.2.3.2 Service Frameworks	35
4.2.4 Names, Description, and Tags	36
4.3 Plan Schema	36
4.3.1 General Nodes	37
4.3.1.1 name	37
4.3.1.2 description	37
4.3.1.3 tags	37
4.3.2 Plan	37
4.3.2.1 camp_version	37
4.3.2.2 origin	37
4.3.2.3 artifacts	38
4.3.2.4 services	38
4.3.3 ArtifactSpecification	38
4.3.3.1 type	38
4.3.3.2 content	38
4.3.3.3 requirements	38
4.3.4 ContentSpecification	39
4.3.5 RequirementSpecification	39
4.3.5.1 type	40
4.3.5.2 fulfillment	40
4.3.6 ServiceSpecification	40
4.3.6.1 id	40
4.3.6.2 href	40
4.3.6.3 characteristics	41
4.3.7 CharacteristicSpecification	41
4.3.7.1 type	41
5 Resources	42
5.1 Required Attributes	42
5.2 Attribute Types	42
5.2.1 Boolean	42
5.2.2 String	42
5.2.3 Number	42
5.2.4 URI	42
5.2.5 Timestamp	42
5.3 CAMP Resource Type Inheritance	42
5.4 camp_resource Resource	43
5.4.1 uri	43
5.4.2 name	43
5.4.3 description	43
5.4.4 tags	43
5.4.5 representation_skew	43

5.4.6 external_management_resource	44
5.4.7 metadata.....	44
5.4.7.1 type_definition.....	44
5.4.7.2 mutable	44
5.4.7.3 consumer_mutable	45
5.5 HTTP Method Support.....	45
5.6 collection Resource	45
5.6.1 collection_type.....	46
5.6.2 total_items	46
5.6.3 items_per_page.....	46
5.6.4 start_index.....	46
5.6.5 items	47
5.7 platform_endpoints Collection	47
5.8 platform_endpoint Resource.....	47
5.8.1 platform.....	48
5.8.2 specification_version	48
5.8.3 backward_compatible_specification_versions	48
5.8.4 implementation_version	48
5.8.5 backward_compatible_implementation_versions.....	49
5.8.6 auth_scheme	49
5.9 platform Resource	49
5.9.1 supported_format_collection	50
5.9.2 extension_collection	50
5.9.3 type_definition_collection	50
5.9.4 platform_endpoints_collection.....	50
5.9.5 specification_version	51
5.9.6 implementation_version	51
5.9.7 assembly_factory	51
5.9.8 service_collection	51
5.9.9 plan_factory.....	51
5.10 assembly_factory Resource	51
5.10.1 parameter_definition_collection	52
5.11 assembly Resource	52
5.11.1 component_collection.....	53
5.11.2 plan.....	53
5.11.3 operation_collection	53
5.11.4 sensor_collection.....	53
5.12 component Resource.....	53
5.12.1 assembly_collection	54
5.12.2 artifact.....	54
5.12.3 service	54
5.12.4 status	54
5.12.5 related_component_collection.....	54
5.12.6 operation_collection	54
5.12.7 sensor_collection.....	55
5.13 service Resource	55

5.13.1	parameter_definition_collection	55
5.13.2	characteristics.....	55
5.14	plan_factory Resource.....	56
5.14.1	parameter_definition_collection	56
5.15	plan Resource.....	56
5.15.1	Advertising Support for the Plan Resource	58
5.16	format Resource	58
5.16.1	mime_type	58
5.16.2	version	59
5.16.3	documentation	59
5.16.4	Required JSON Format Resource	59
5.17	type_definition Resource	59
5.17.1	items	60
5.17.2	documentation	60
5.17.3	inherits_from_collection.....	60
5.18	attribute_definition Resource	60
5.18.1	documentation	60
5.18.2	attribute_type.....	61
5.18.3	required	61
5.19	parameter_definition Resource	61
5.19.1	parameter_type	61
5.19.2	required	61
5.19.3	default_value	62
5.19.4	parameter_extension.....	62
5.20	operation Resource	62
5.20.1	name.....	62
5.20.2	documentation	63
5.20.3	target_resource	63
5.20.4	parameter_definition_collection	63
5.21	sensor Resource.....	63
5.21.1	documentation	63
5.21.2	target_resource	64
5.21.3	sensor_type	64
5.21.4	value	64
5.21.5	units	64
5.21.6	timestamp	64
5.21.7	operation_collection	64
6	Protocol	65
6.1	Transfer Protocol	65
6.2	URI Space.....	65
6.3	Media Types	65
6.3.1	Required Formats.....	65
6.3.1.1	Duplicate Keys in JSON Objects.....	65
6.3.2	Supported Formats.....	65
6.4	Request Headers	66

6.5	Query String Parameters	66
6.6	POST Body Parameters	66
6.6.1	Parameter Handling	67
6.7	Response Headers	67
6.8	HTTP Status Codes	67
7	Operations	68
7.1	Deploying an Application	68
7.1.1	Deploying an Application by Reference	68
7.1.2	Deploying an Application by Value	69
7.1.2.1	Deploying an Application by Value Using MIME	69
7.1.2.2	Deploying an Application by Value Without MIME	70
7.2	Registering a Plan	70
7.2.1	Registering a Plan by Reference	70
7.2.2	Registering a Plan by Value	71
7.2.2.1	Registering a Plan by Value Using MIME	71
7.2.2.2	Registering a Plan by Value Without MIME	72
7.3	Retrieving Resources	72
7.3.1	Selecting Attributes on a GET	72
7.3.2	Selecting Attributes of the Resources in a Collection on a GET	72
7.3.2.1	Uniqueness of Resources in a Collection	73
7.3.3	Sorting Elements of a Collection	74
7.3.3.1	Sort Ordering	75
7.3.4	Pagination of a Collection	75
7.3.5	Retrieving Index of an Item in a Collection	76
7.4	Updating Resources	76
7.4.1	Updating with PUT	76
7.4.1.1	Partial Updates with PUT	77
7.4.2	Updating with JSON Patch	77
7.4.3	Mutability of Resource Attributes	77
7.5	Stopping an Application Instance	77
8	Extensions	78
8.1	Unique Name Requirement	78
8.2	extension Resource	79
8.2.1	version	79
8.2.2	documentation	80
8.3	Extending Existing Resources	80
9	Conformance	81
9.1	CAMP Provider	81
9.2	CAMP Consumer	81
9.3	Platform Deployment Package	81
9.4	Plan	81
Appendix A.	Acknowledgments	82
Appendix B.	Normative Statements	83
B.1	Mandatory Statements	83
B.2	Non-Mandatory Statements	90
Appendix C.	Example Version Scheme	94

Appendix D. Revision History 95

1 Introduction

1.0 IPR Policy

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

1.1 Overview

Platform as a Service (PaaS) is a term that refers to a type of cloud computing in which the service provider offers customers/consumers access to one or more instances of a running application computing platform or application service stack. NIST defines PaaS [\[SP800-145\]](#) as a “service model” with the following characteristics:

The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.

There are multiple commercial PaaS offerings in existence using languages such as Java, Python and Ruby and frameworks such as Spring, Django and Rails. Although these offerings differ in such aspects as programming languages, application frameworks, etc., there are inherent similarities in the way they manage the lifecycle of the applications that are targeted for, and deployed upon them. *The core proposition of this specification is that these similarities can be leveraged to produce a generic application and platform management API that is language, framework, and platform neutral.*

For PaaS consumers this management API would have the following benefits:

- “Portability between clouds” is emerging as one of the primary concerns of cloud computing. By standardizing the management API for the use cases around deploying, stopping, starting, and updating applications, this specification increases consumers’ ability to port their applications between PaaS offerings.
- It is likely that implementations of this specification will appear as plugins for application development environments (ADEs) and application management systems. Past experience has shown that, over time, such generic implementations are likely to receive more attention and be of higher quality than the implementations written for solitary, proprietary application management interfaces.

For PaaS providers this management API would have the following benefits:

- Because the strength and features of a PaaS offering’s application management API are unlikely to be perceived as key differentiators from other PaaS offerings, the existence of a consensus management API allows providers to leverage the experience and insight of the specification’s contributors and invest their design resources in other, more valuable areas.
- By increasing the portability of applications between PaaS offerings, this management API helps “grow the pie” of the PaaS marketplace by addressing one of the key pain points for PaaS consumers.

1.2 Purpose

This document defines the artifacts and APIs that need to be offered by a Platform as a Service (PaaS) cloud to manage the building, running, administration, monitoring and patching of applications in the cloud. Its purpose is to enable interoperability among self-service interfaces to PaaS clouds by defining

artifacts and formats that can be used with any conforming cloud and enable independent vendors to create tools and services that interact with any conforming cloud using the defined interfaces. Cloud vendors can use these interfaces to develop new PaaS offerings that will interact with independently developed tools and components.

The following is a non-exhaustive list of the use cases which are supported by this specification.

- Building and packaging an application in a local Application Development Environment (ADE)
- Building an application in an ADE running in the cloud
- Importing a Platform Deployment Package into the cloud
- Uploading application artifacts into the cloud
- Run, stop, suspend, snapshot, and patch an application

1.3 Example (non-normative)

This example illustrates a scenario in which the application administrator wants to run and monitor an application.

The administrator does this by deploying the application, in the form of a *Platform Deployment Package*, to the platform. This is done by sending an HTTP POST request to the URL of the *assembly_factory resource* as shown below, where `"/my_paas/assembly_factory"` is this URL and `"/my_paas/pkgs/1"` is the location of the application package.

```
POST /my_paas/assembly_factory HTTP/1.1
Host: example.org
Content-Type: application/json
Content-Length: ...

{
  "pdp_uri": "/my_paas/pkgs/1"
}
```

On receiving such a request the platform unpacks the package, parses and validates the Plan file, resolves the service dependencies described by that Plan, and starts the application. On successful start the platform creates a new resource representing the running application and provides the URL of that resource `"/my_paas/apps/1"` in the response as shown below.

```
HTTP/1.1 201 Created
Location: http://example.org/my_paas/apps/1
Content-Type: ...
Content-Length: ...

...
```

The administrator can now monitor the running application by sending an HTTP GET request to the resource that represents the running application, which was obtained in the previous step (`"/my_paas/apps/1"`).

```
GET /my_paas/apps/1 HTTP/1.1
Host: example.org
```

The response contains the JSON representation of the running application as shown below.

```
HTTP/1.1 200 OK
Content-Type: application/json
Content-Length: ...

{
  "uri": "http://example.org/my_paas/apps/1",
  "name": "Hello Cloud App",
  "description": "Hello Cloud Application Running in a PaaS Env",
  "component_collection": "http://example.org/collections/aad7",
  "metadata": {
    "type_definition": "http://example.org/my_paas/types/assembly",
    "mutable": ["/name", "/description"],
    "consumer_mutable": ["/description"]
  }
}
```

1.4 Non-Goals

The interfaces exposed by the components and services in a PaaS system can be broadly split into two categories; functional interfaces and management interfaces. Functional interfaces are those that involve the specific utility provided by that component. For example, the interface used to submit a message to a message queuing service is as a functional interface. Management interfaces are those that deal with the administration of components. For example, the interface used to deploy and start an application on the platform is a management interface.

The specification of functional interfaces is out of scope for this document.

1.5 Actors

There are many actors for a PaaS environment. For the purposes of this specification we identify the following actors:

Application Developer: The person that builds and tests an application and presents the developed artifacts for deployment.

Application Administrator: The person that deploys applications and manages the application throughout its life-cycle.

Together these two actors make up the consumers of the management API described in this specification. This specification is intended mainly for Application Administrators, though it does constraint the artifacts that an Application Developer presents for deployment.

Platform Administrator: The person that manages the platform. This specification describes some of the functions of a Platform Administrator, though most of the functions of this actor are outside its scope.

Application End-User: A user of an application deployed on the platform. The interactions of the Application end-user and the application are outside the scope of this specification.

Extension Developer: The person who creates new Extensions for Platforms.

1.6 Terminology

1.6.1 Term Definitions

1.6.1.1 CAMP Provider (Provider)

A CAMP Provider (Provider) is an implementation of the service aspects of this specification.

1.6.1.2 CAMP Consumer (Consumer)

A CAMP Consumer (Consumer) is an implementation of the client aspects of this specification.

1.6.2 Keywords, Conventions, and Normative Text

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 [RFC2119].

Upper case versions of the RFC 2119 keywords are used in this document for emphasis. All text except examples, unless otherwise labeled, is normative. Normative statements that use these keywords have been highlighted as per this sentence. Each such statement has been given a unique tag in the following manner: [EX-00]. For convenience these statements have been tabulated and cross-indexed by their tags and appear in Appendix B. All examples, figures and notes in this document are informative. Unless marked otherwise text in this specification is normative.

See Section 9, “Conformance”, for details on Conformance to this specification.

1.7 Notational Conventions

The JSON and YAML descriptions that depict the representation of resources and the structure of Plans use a pseudo-schema notation with the following conventions:

- Characters are appended to items to indicate cardinality:
 - “?” (0 or 1)
 - “*” (0 or more)
 - “+” (1 or more)

Absent any indication of cardinality, the default cardinality of an element is “exactly 1”. The scope of these operators is the entire line on which they appear.

- Vertical bars, “|”, denote choice. For example, “a | b” means a choice between “a” and “b”.
- Parentheses, “(“ and “)”, are used to indicate the scope of the “|” operator.
- An expression in italics indicates a value whose type is indicated by the italicized expression. For example, “foo: *String*” indicates that the value of “foo” will be a String.
- Square brackets, “[]”, indicate an array of the type indicated by the expression preceding it. For example, “foo: *String*[]” indicates that the value of “foo” will be an array of Strings.
- An expression surrounded in angle brackets, “<” and “>”, indicates a value whose type is indicated either by some of field in the object or by other context information. For example, “foo: <sensor_type>” indicates that the type of the value of “foo” is provided by the value of the “sensor_type” attribute.

Note that the information presented in pseudo-schema is not normative and is intended as a condensed guide for ease of understanding by the reader. In the event of a conflict (due to a typo or other editorial error) the text takes precedence over the pseudo-schema.

1.8 Specification Version

Each version of a CAMP specification is identified by a unique string termed the “Specification Version String”. The Specification Version String for this specification is “CAMP 1.2”.

1.8.1 Backwards Compatibility

This version of the CAMP specification is not backwards compatible with any previous version of the CAMP specification.

1.9 Normative References

- [ISO 8601:2004] International Organization for Standardization, Geneva, Switzerland, “Data elements and interchange formats -- Information interchange - - Representation of dates and times”, March 2008. http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=40874

- [OVF] DMTF DSP0243, "Open Virtualization Format Specification 2.0.1", http://www.dmtf.org/sites/default/files/standards/documents/DSP0243_2.0.1.pdf
- [RFC1952] Deutsch, P., "GZIP file format specification version 4.3", RFC 1952, May 1996. <http://www.ietf.org/rfc/rfc1952.txt>
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997. <http://www.ietf.org/rfc/rfc2119.txt>.
- [RFC2388] Masinter, L., "Returning Values from Forms: multipart/form-data", RFC 2388, August 1998. <http://www.ietf.org/rfc/rfc2388.txt>
- [RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999. <http://www.ietf.org/rfc/rfc2616.txt>
- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, May 2000. <http://www.ietf.org/rfc/rfc2818.txt>
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005. <http://www.ietf.org/rfc/rfc3986.txt>
- [RFC4346] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.1", RFC 4346, April 2006. <http://www.ietf.org/rfc/rfc4346.txt>
- [RFC4627] Crockford, D., "The application/json Media Type for JavaScript Object Notation (JSON)", RFC 4627, July 2006. <http://www.ietf.org/rfc/rfc4627.txt>
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008. <http://www.ietf.org/rfc/rfc5246.txt>
- [RFC5789] Dusseault, L. and J. Snell, "PATCH Method for HTTP", RFC 5789, March 2010. <http://www.ietf.org/rfc/rfc5789.txt>
- [RFC6901] Bryan P., K. Zyp, M. Nottingham, "JavaScript Object Notation (JSON) Pointer", RFC 6901, April 2013. <http://www.ietf.org/rfc/rfc6901.txt>
- [RFC6902] Bryan, P., Ed., and M. Nottingham, Ed., "JavaScript Object Notation (JSON) Patch", RFC 6902, April 2013. <http://www.ietf.org/rfc/rfc6902.txt>
- [SHA256] FIPS PUB 180-4, Federal Information Processing Standards Publication, "Secure Hash Standard (SHS) (FIPS PUB 180-4)", 6.2, SHA-256. <http://csrc.nist.gov/publications/fips/fips180-4/fips-180-4.pdf>
- [TAR] ISO/IEC/IEEE 9945:2009(E), International Standard - Information technology Portable Operating System Interface (POSIX) Base Specifications, Issue 7. <http://ieeexplore.ieee.org/document/5393893/>
- [UTS10] Unicode® Technical Standard #10, "Unicode Collation Algorithm", June 16, 2017. <http://www.unicode.org/reports/tr10/tr10-36.html>
- [YAML 1.1] Oren Ben-Kiki, Clark Evans, Brian Ingerson, "YAML Ain't Markup Language (YAML) Version 1.1, 2005-01-18". <http://yaml.org/spec/1.1/>. Also archived at <http://xml.coverpages.org/yaml-spec-v1.1-archive-copy.html>
- [ZIP] ZIP File Format Specification, <https://pkware.cachefly.net/webdocs/APPNOTE/APPNOTE-6.3.0.TXT>

1.10 Non-Normative References

- [Git] The Software Freedom Conservancy, "Git, the fast version control system", March 2012. <https://git-scm.com/>
- [Keystone] OpenStack Foundation, "Identity Service API Reference v3.8", September 2017. <https://developer.openstack.org/api-ref/identity/v3/index.html>
- [RFC2617] Franks, J., Hallam-Baker, P., Hostetler, J., Lawrence, S., Leach, P., Luotonen, A., and L. Stewart, "HTTP Authentication: Basic and Digest Access Authentication", RFC 2617, June 1999. <http://www.ietf.org/rfc/rfc2617.txt>
- [RFC6749] Hardt, D., Ed., "The OAuth 2.0 Authorization Framework", RFC 6749, October 2012. <http://www.ietf.org/rfc/rfc6749.txt>

[SP800-145]

Peter Mell, Timothy Grance, "The NIST Definition of Cloud Computing", Special Publication 800-145, September 2011.

<http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf>

[POM-Syntax]

Tim O'Brien, Manfred Moser, John Casey, Brian Fox, Jason Van Zyl, Eric Redmond Larry Shatzer, "Maven: The Complete Reference", 2008-2014, Section 3.3 POM Syntax. <https://www.sonatype.com/ebooks>

2 Concepts and Types

This section is informative.

2.1 Introduction

This specification defines the self-service management API that a Platform as a Service offering presents to the consumer of the platform. The API is the interface into a platform implementation layer that controls the deployment of applications and their use of the platform.

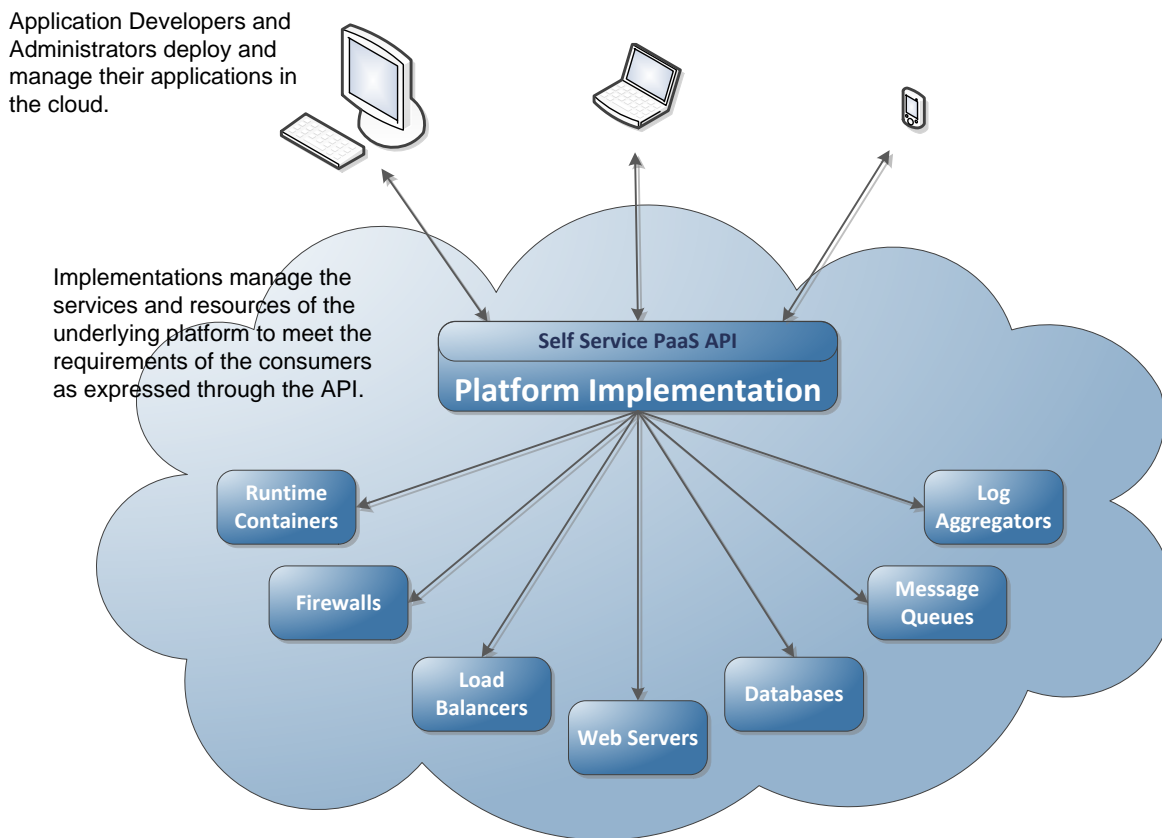


Figure 2-1: Typical PaaS Architecture

The figure above shows a typical architecture of a Platform as a Service cloud. The platform implementation is a management client of the underlying resources that transforms (through policies) the application requirements expressed by the Application Administrator into provisioning and other operations on those resources. The Platform Administrator manages the underlying hardware, storage, networks, and software services that make up the platform through existing administrative interfaces. Thus the Application Administrator is able to concentrate on their application and its deployment environment rather than having to be a systems administrator, database administrator and middleware administrator as well (as is the case with IaaS).

The goal of the management interface is to provide the PaaS consumer with a model that is as simple as possible, and yet still provides the essential elements that give them control over the deployment, execution, administration and metering of their application and its deployment environment.

2.2 Resources

The CAMP API is made up of resources in a REST protocol. The resources represent elements of the underlying system. The protocol enables interaction with the resources. The following are the main resources in the API:

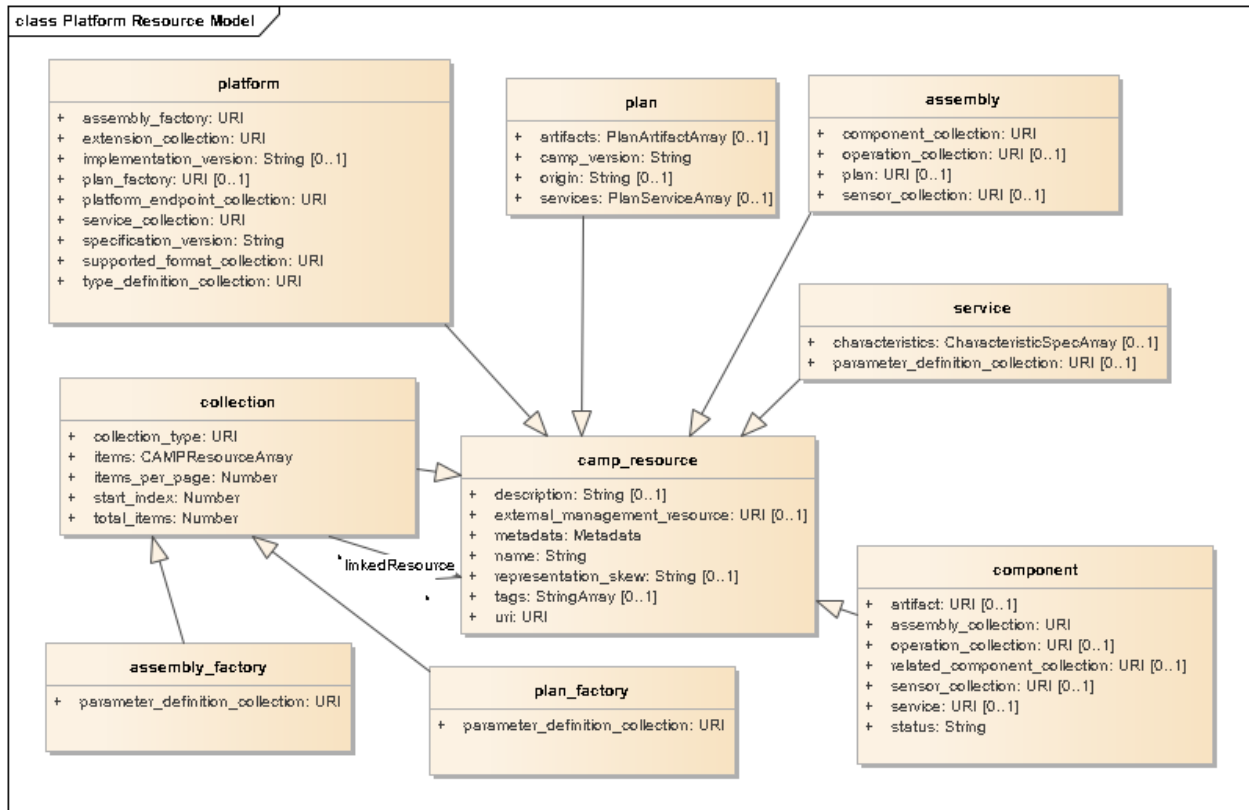


Figure 2-2: CAMP Resources as UML Classes

Figure 2-2 is a UML Class Diagram showing the CAMP resources as UML classes. All CAMP resources share a set of common attributes which they inherit from the *camp_resource* parent class.

Each attribute shown in these UML class diagrams has a CAMP common attribute type. The '+' symbol before each attribute name in the boxes indicates that the attribute access is public (i.e. available through the API). Non-mandatory resource attributes are indicated using the [0..1] UML multiplicity tag.

2.2.1 Platform

The *platform resource* is the primary view of the platform and what is running on it. The *platform resource* references collections of resources that represent the services provided by the platform (as *Services*), the applications running on this platform (as *assembly resources*), as well as collections of metadata resources that describe the resources supported by the platform as well as any extensions that the Provider has implemented. The *platform resource* also determines the scope of access for sharing amongst multiple applications.

2.2.2 Assemblies

An *assembly resource* represents running applications. Operations on an *assembly resource* affect the components and elements of that application.

2.2.3 Components

An *assembly resource* is comprised of one or more *component resources*. A *component resource* represents a discrete and, in most cases, dynamic element of an application such as such as a deployed Ruby gem, a database instance, or a set of entries in a LDAP directory. A *component resource* can be related to other *component resources* through producer/consumer or other kinds of relationships.

2.2.4 Plans

A Plan is meta-data that provides a description of the artifacts that make up an application, the services that are required to execute or utilize those artifacts, and the relationship of the artifacts to those services. Plans can be expressed in two forms; either as a YAML file or, optionally, as a *plan resource*. The Artifacts described in a Plan represent discrete, static elements of an application such as a Ruby gem file, an SQL script, or a PKCS #12 file.

2.2.5 Services

A *service resource* represents a blueprint for creating *component resources* that utilize or embody a platform-provided service in some way. For example, a Service may represent the platform's ability to create a message queue for use by an application.

2.2.6 Collection and Factory Resources

A resource may refer to a set of other resources through the use of a *collection resource*. For example, an *assembly resource* references the set of components that make up that Assembly through a reference to a collection of *component resources*. The type of the resources in a collection is homogenous. All of the resources in a collection are instances of a given type of resource, or instances of a sub-type of that type.

The *assembly_factory* and *plan_factory resources* are sub-types of the generic *collection resource*. In addition to acting as collections, these resources can be used to create new instances of the resources of which they are a collection (i.e. the *assembly_factory resource* can be used to create new *assembly resources* and the *plan_factory resource* can be used to create new *plan resources*).

2.2.7 Operations and Sensors

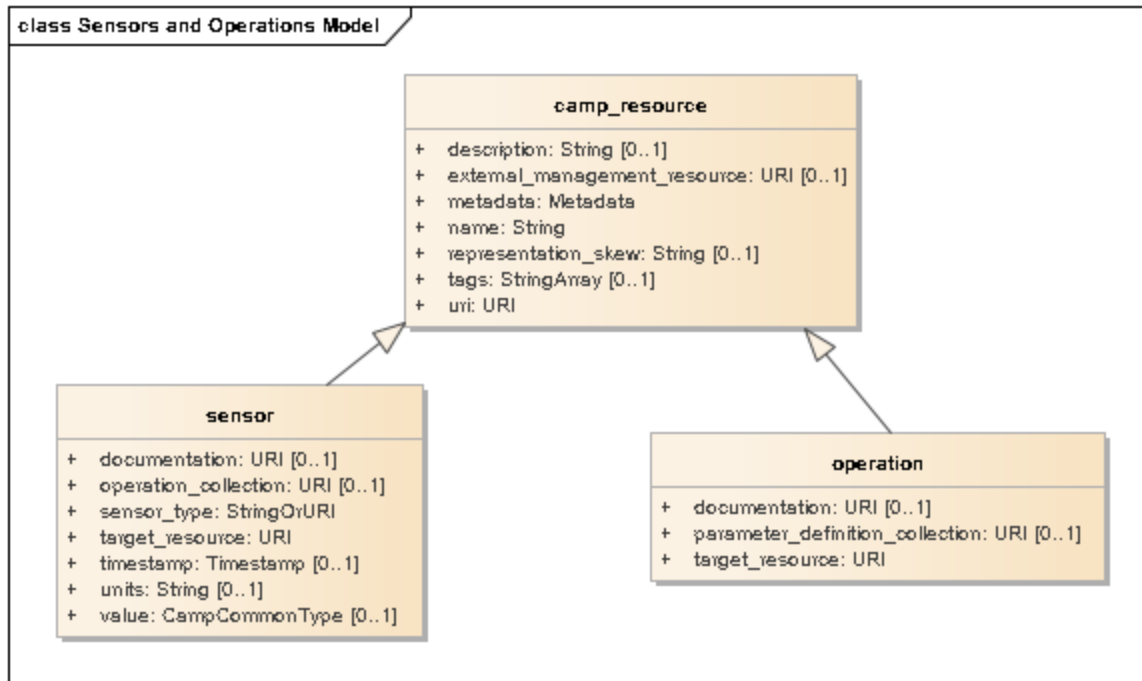


Figure 2-3: Operations and Sensors

Figure 2-3 is a UML class diagram showing the attributes of the *operation resources* and *sensor resources*.

Operations and Sensors provide a way of interacting with an application through the CAMP API. An *operation resource* represents actions that can be taken on a resource, and *sensor resources* represent dynamic data about resources, such as metrics or state. A *sensor resource* is useful for exposing data that changes rapidly, or that might need to be fetched from a secondary system. A *sensor resource* can also offer Operations to allow resetting metrics, or adjusting frequency collection settings.

Multiple *operation resources* and *sensor resources* can be exposed both on *assembly resources* and *component resources*. Operations are also known as effectors. The combination of Operations and Sensors enables ongoing management. This can include automation techniques such as using policies, event-condition-action paradigms, or autonomic control. A Consumer can use the REST API to perform such management. A Provider can also use them. For example, a *component resource* could be offered that allows for “autoscaling” capacity based on the volume of work an application performs.

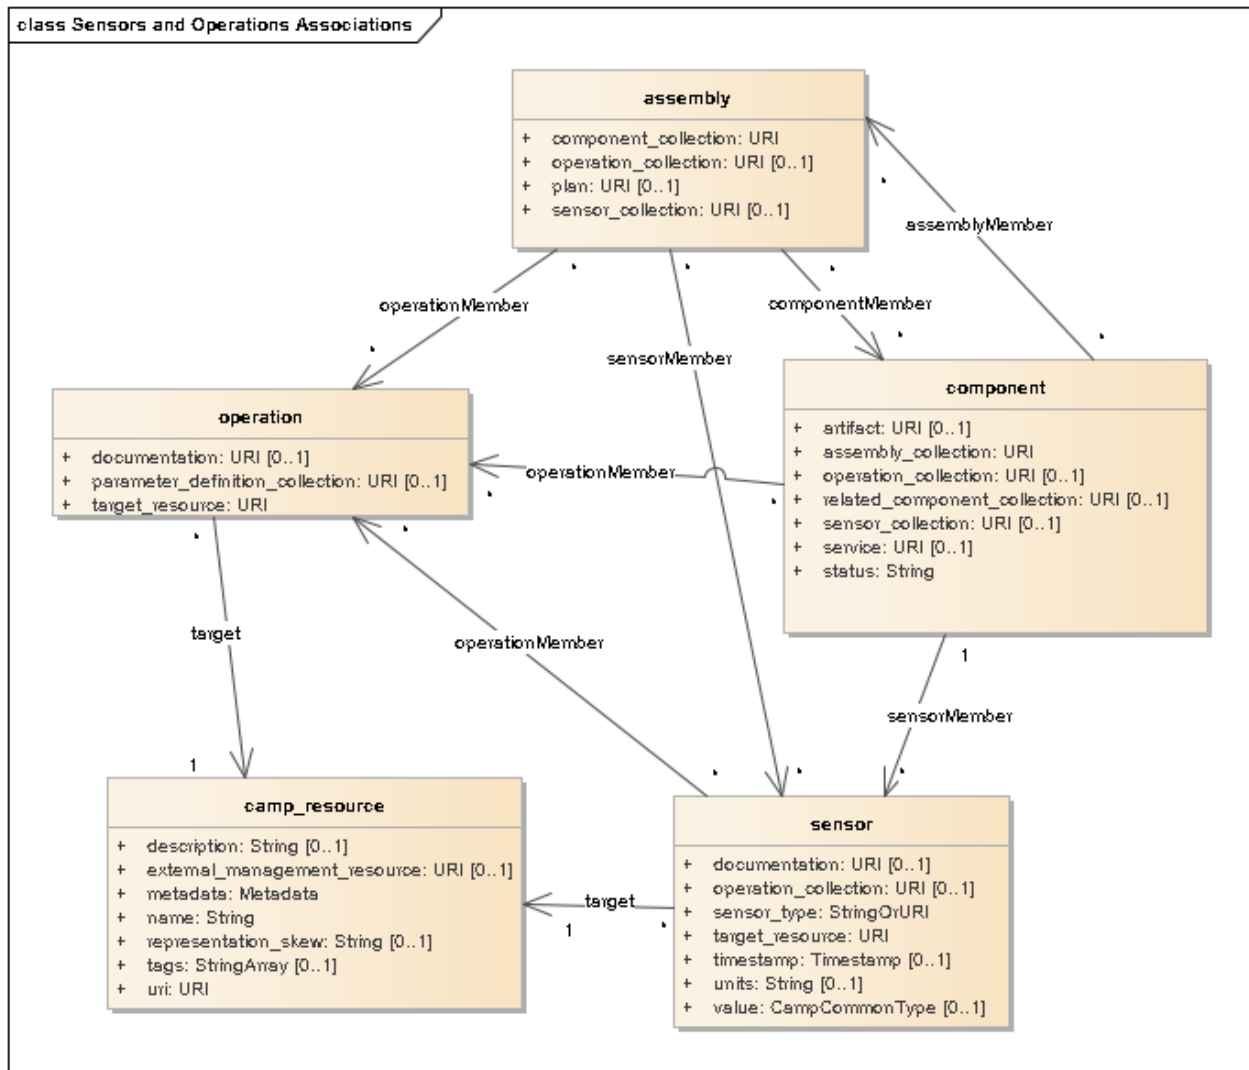


Figure 2-4: Operations and Sensors Associations

Figure 2-4 is a UML class diagram showing *operation resources*, *sensor resources*, and the other CAMP resources that they are associated with.

In Figure 2-4 the notation `<resourceType>Member` is used for association names to indicate that the associated resource type may be a member of a collection linked through a *collection resource* pointed at by a collection attribute in the "from" resource class (i.e. the *collection resource* is not shown in the diagram).

2.2.8 Resource Relationships

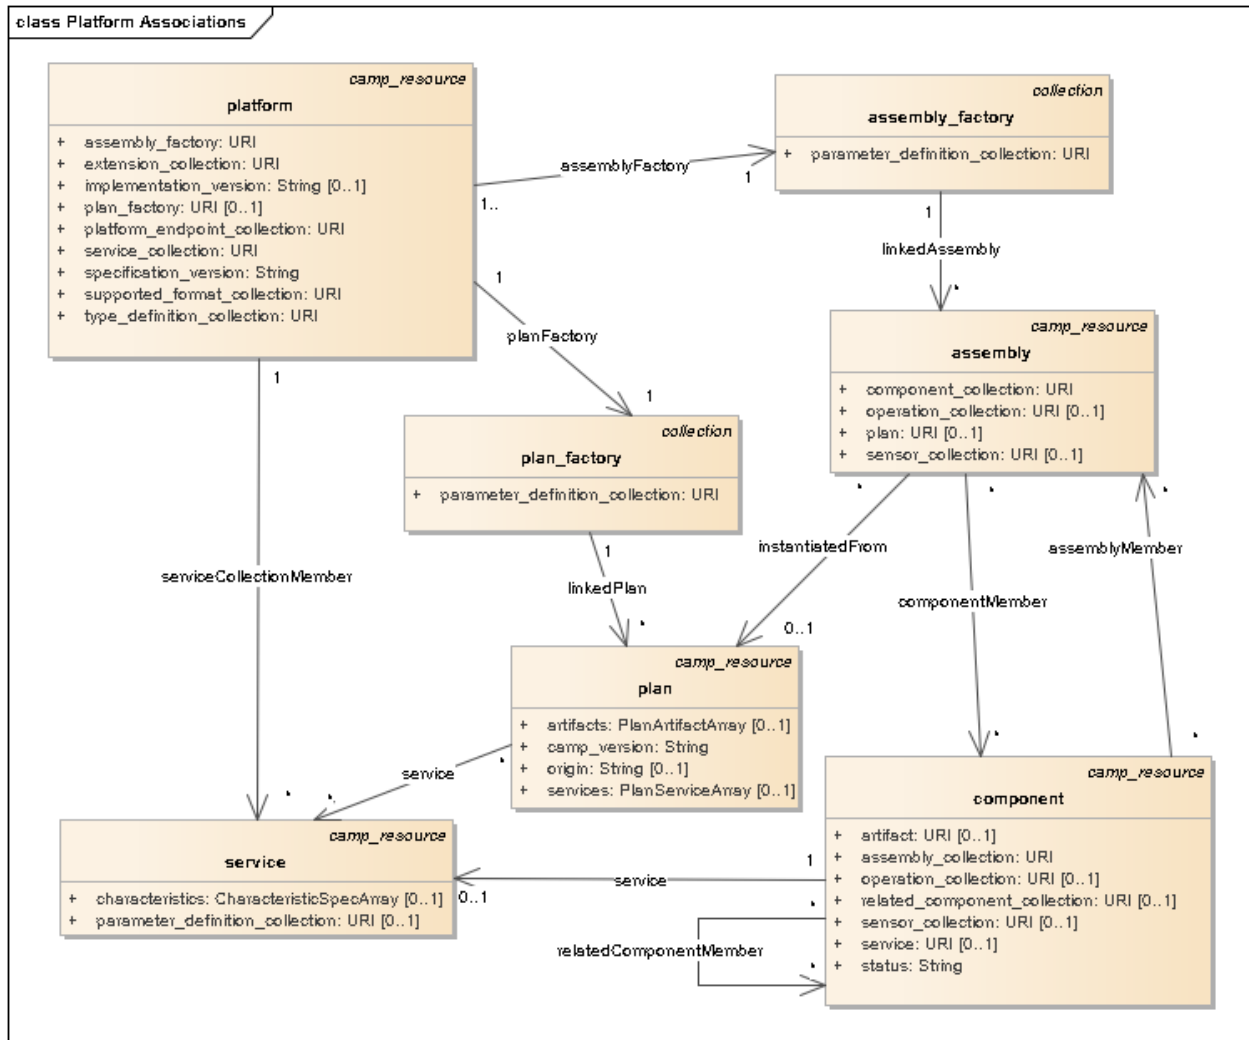


Figure 2-5: Platform and Resource Relationships

Figure 2-5 shows the relationships between Platform Resources using a UML class diagram.

Associations which are visible through pointer attributes in resources (i.e. URIs) are shown using UML named associations with navigation arrows.

The notation <resourceType>Member is used for association names to indicate that the associated resource type may be a member of a collection linked through a *collection resource* pointed at by a collection attribute URI in the "from" resource class (i.e. the *collection resource* is not shown in the diagram).

2.3 Deployment

A *Platform Deployment Package* (PDP) is an archive containing a Plan file together with application content files such as web archives, database schemas, scripts, source code, localization bundles, and icons; and metadata files such as manifests, checksums, signatures, and certificates. It can be used to move an Application and its Components from Platform to Platform, or between an Application Development Environment and a Platform.

In the simplest case (an example of which is provided Section 1.3, "Example"), a PDP or a Plan file can be used to create an *assembly resource* by transmitting an HTTP POST request containing either the PDP or the Plan file to the *assembly_factory resource*.




Figure 2-6: Deploying an Application

On platforms that choose to support Plans, a CAMP Consumer can create a *plan resource* by uploading either a PDP or a Plan file to the *plan_factory resource* using an HTTP POST request. An *assembly resource* can then be created from the *plan resource* by including a reference to the *plan resource* in an HTTP POST request to the *assembly_factory resource*.

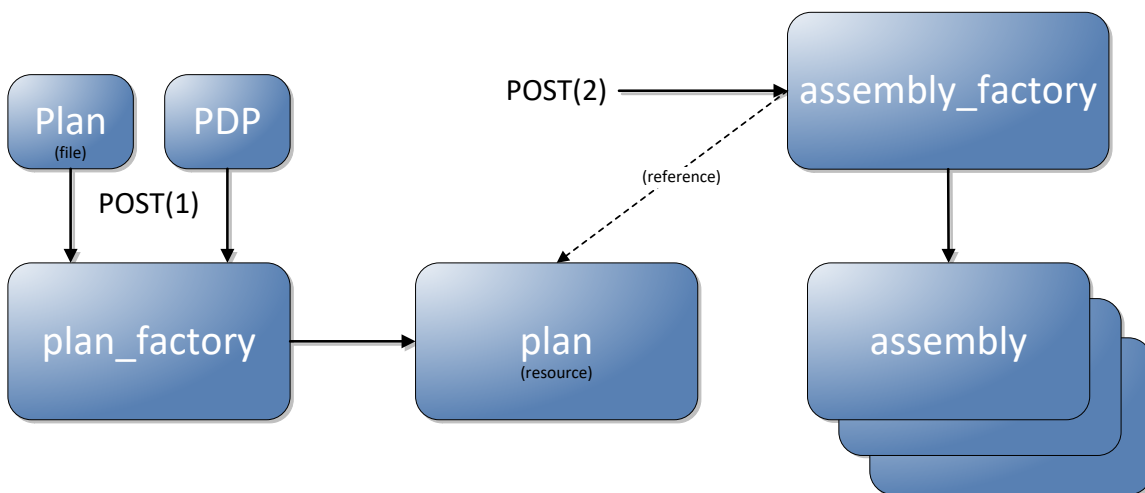


Figure 2-7: Instantiating an Application from a plan resource

In Figure 2-7 the POST(1) request creates a *plan resource* by uploading either a PDP or a Plan file to the *plan_factory resource*. The POST(2) request to the *assembly_factory resource* creates an *assembly resource*. Multiple *assembly resources* can be created from a single *plan resource* by submitting multiple HTTP POST requests.

2.4 Versions and Extensions

This specification supports multiple endpoints and versions, and extensions. All of these are represented in the resource model so they can be discovered by CAMP Consumers. The resources enabling discovery are shown in Figure 2-8, and their relationships are shown in Figure 2-9.

In these figures, the notation <resourceType>Member is used for association names to indicate that the associated resource type may be a member of a collection linked through a *collection resource* pointed at by a collection attribute URI in the "from" resource class (i.e. the *collection resource* is not shown in the diagram).




Figure 2-8: Platform Endpoint and Metadata Resources

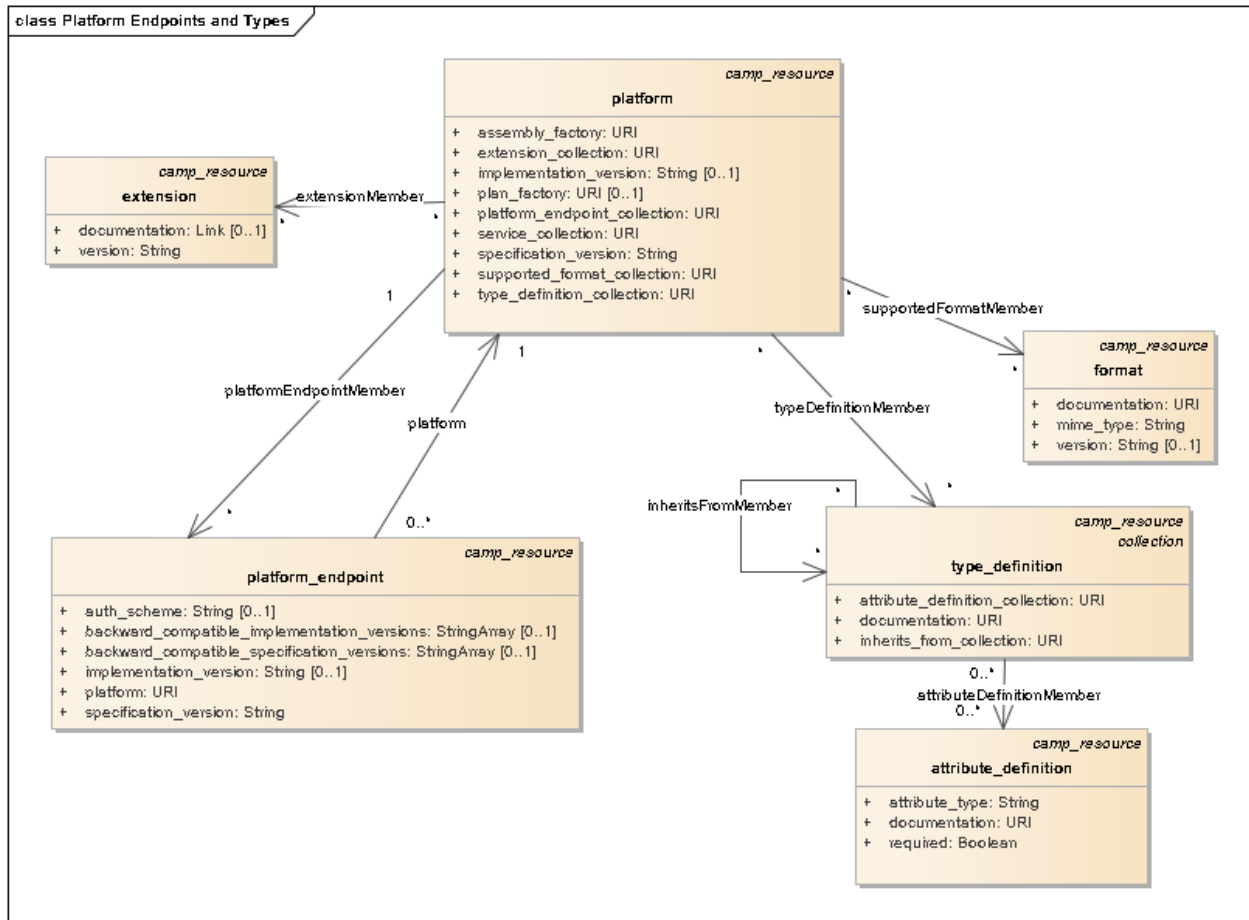


Figure 2-9: Platform Endpoint and Extension Resource Relationships

2.5 Parameters

Parameters can be defined on the *assembly_factory resource*, *service resources*, and, if supported, *plan_factory resource*. Parameters affect the resources that are generated from these resources. Figure 2-10 illustrates the relationships between these resources and the resources used to represent Parameters.

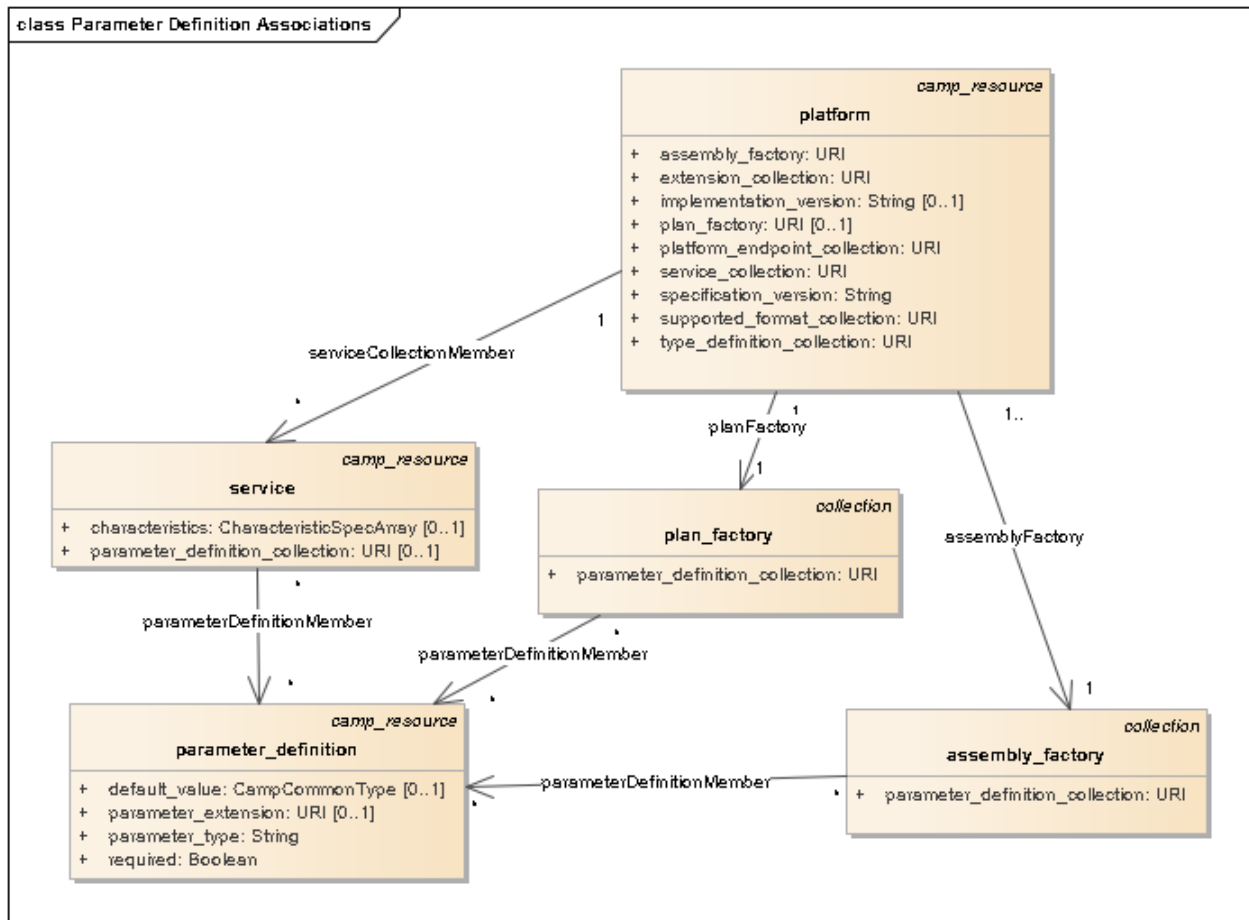


Figure 2-10: Parameter Definition Relationships

2.6 CAMP Common Attribute Types

Many of the attributes in the UML class diagrams have one of the CAMP common attribute types, specified in Section 5.2, “Attribute Types”. Figure 2-11 is a UML diagram showing the common data types as UML Data Types, which are used for the Types of these Resource Attribute definitions.

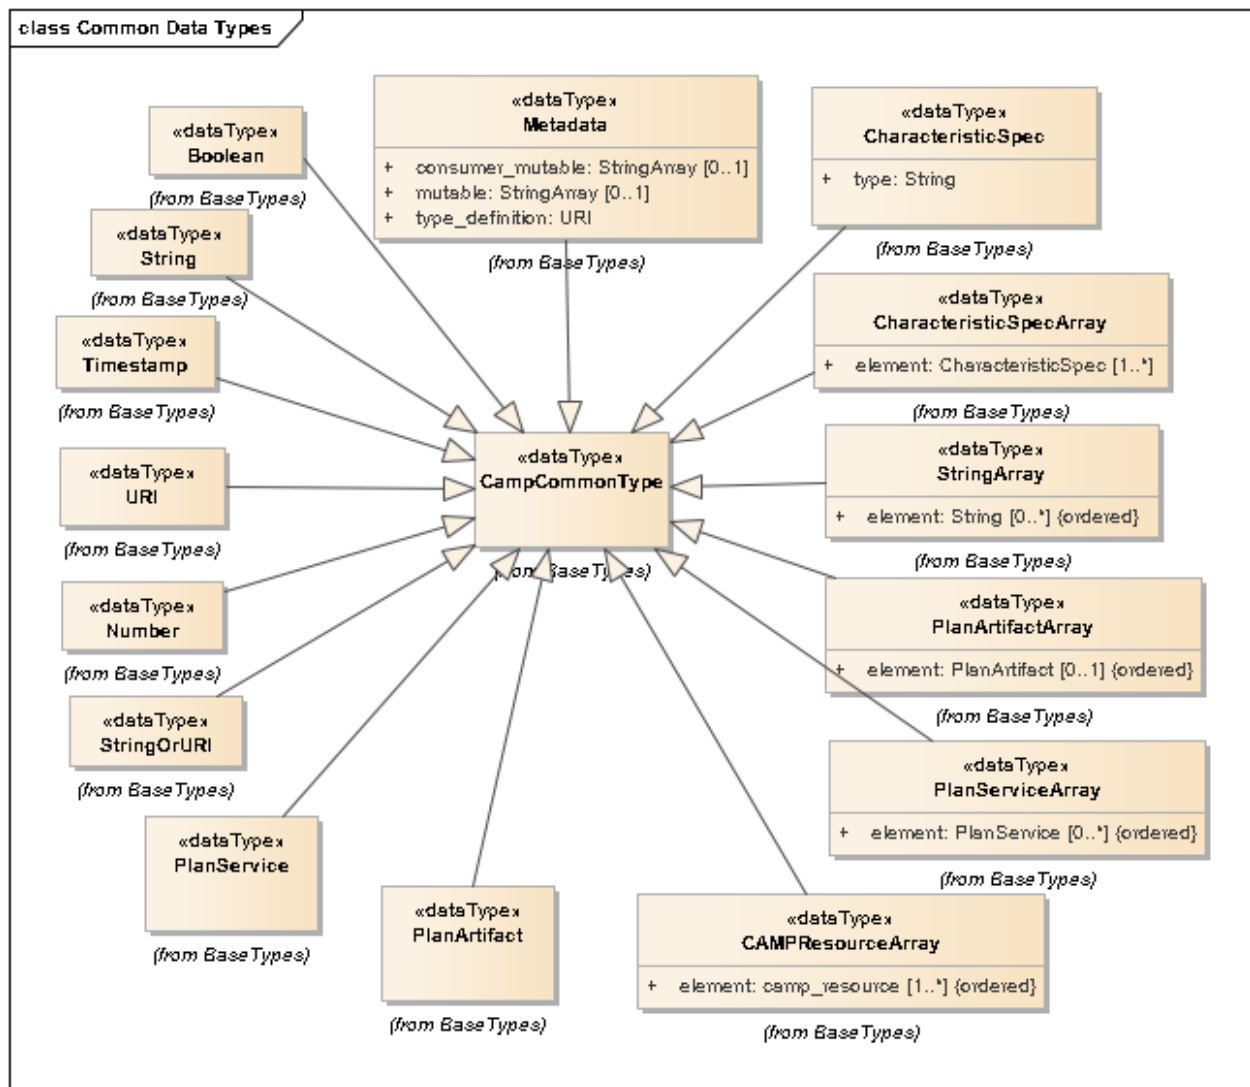


Figure 2-11: CAMP Common Base Types for Resource Attribute Definitions

Multi-valued member attributes are used to model the elements of the StringArray. This is done for modeling purposes only; the attribute name "element" does not appear in the JSON serialization for these common types.

The array types have their elements tagged as ordered.

2.7 Representation Skew

There can be situations in which the information in the resources provided by the CAMP API is not a complete or accurate representation of the state of the underlying implementation. For example, while generating a new instance of an application, a CAMP server might be asked to provide a representation of a Component that corresponds to a dataset that is in the process of being loaded onto a database. While the CAMP server might not be able to provide all of the information about this Component, it would be inaccurate to say that the Component does not exist; it exists but in an intermediate state. It is expected that these sorts of situations will be the exception and that, during the majority of its existence, a CAMP resource will be in synch with the state of its underlying platform implementation.

The significance of this skew is the manner in which it affects the Consumer's interactions with, and expectations about, the resource. In the above example, while the Consumer cannot make any changes to the Component until it has reached a steady state, the Consumer can expect that the resource will reach this state in the near future. There are other situations in which, through some sort of error, the

CAMP API cannot tell when or if the information in the resource will ever be synchronized with the underlying implementation.

Details on how this skew is exposed in the CAMP API are provided in Section 5.4.5, "[representation_skew](#)".

2.8 External Management Resources

The information provided by the resources defined in the CAMP API is only one view of the implementation of the underlying platform. There are, of course, other ways of viewing and interacting with the entities that are modeled by CAMP's resources. For example, CAMP may represent an instance of a database as a *component resource*. This database instance may also be viewed and managed through a database management application. To enable the integration of CAMP-based management tools with other management systems all CAMP resources have an optional URI attribute that references an external management resource. In the above example, the *component resource* that represented the database instance could have an `external_management_resource` attribute that, when de-referenced, would bring the Consumer to a database management system with that database instance already contextualized (i.e. "selected" in whatever manner the UI did such things).

3 Application Management Lifecycle

This section is informative. The figures in this section are UML object instance diagrams, which represent related Resources at various stages of Platform Resource lifecycle. For simplification, attributes for these resources are not shown. For a comprehensive list of attributes for resources see Section 5, “Resources”.

Instances in these diagrams are indicated by boxes, with an underlined “*object-name: Class*” label.

Relationships visible through the API are shown using associations with navigation arrows.

Implementation specific relationships are indicated using the association end notation, without navigation arrows.

3.1 Initial Platform Resources

The CAMP model includes the resources below when no *assembly resources* or *plan resources* have been created. Note that the support of the *plan resource* collection and *plan resources* is optional.

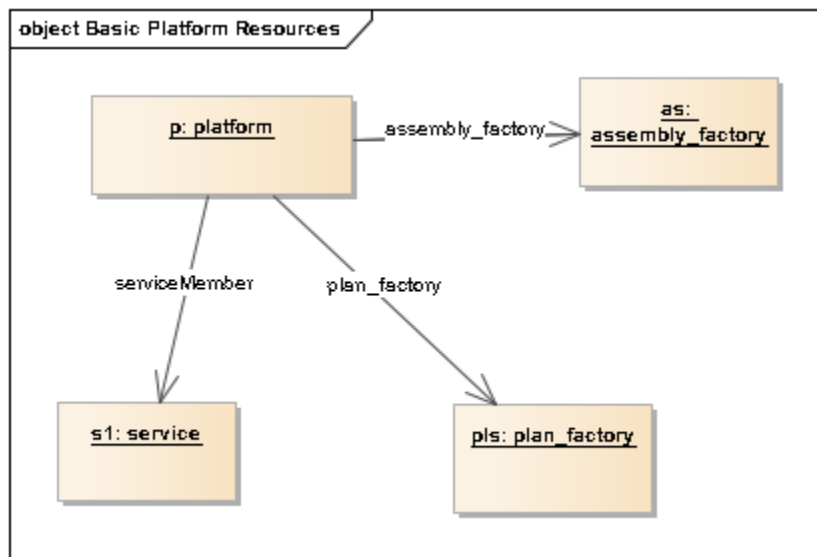


Figure 3-1: Initial Platform Resources

When the Application Administrator first accesses a new account a Platform will have a number of resources visible through the API. The *platform resource* is used to find the other resources in this diagram. The various *service resources* allow for discovery of all the platform services that are available along with value ranges for each service’s attributes.

3.2 Creating an Assembly from a PDP or Plan File

A CAMP Consumer can create a new *assembly resource* by uploading either a PDP or a Plan file to the *assembly_factory resource* using an HTTP POST request (see Section 2.3, “Deployment”). The loaded assembly model might then appear as follows (for simplification, the instantiated component resources are not shown in Figure 3-2):

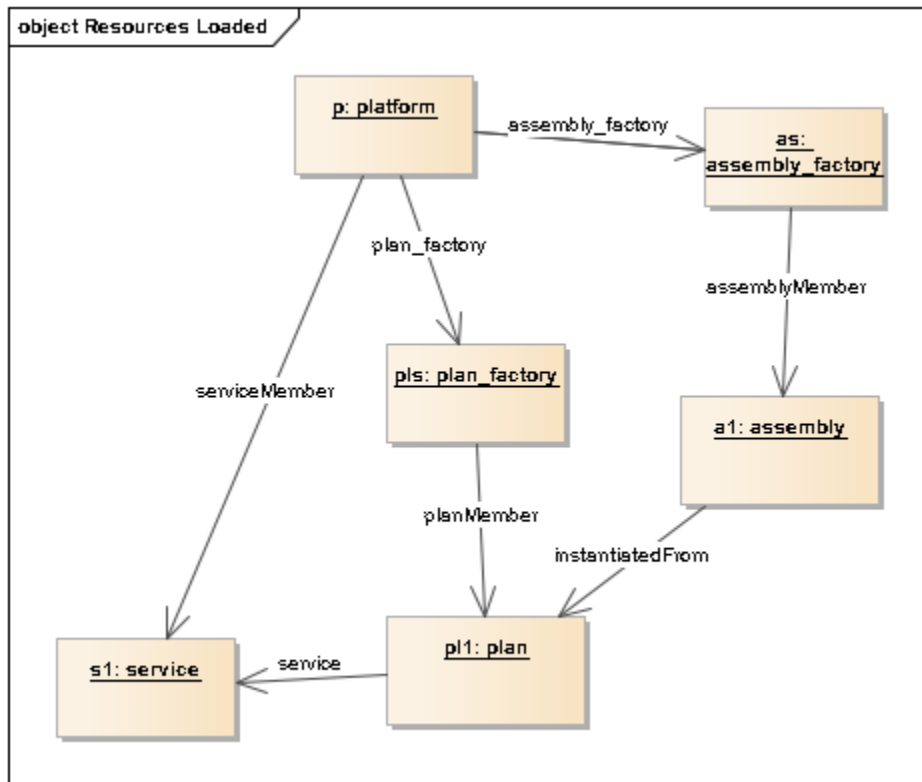


Figure 3-2: Loaded Assembly Resource

If any of its requirements are not resolved, a PDP or Plan file could require modification before it can be used to create *assembly resources*.

3.3 Creating an Assembly from a plan resource

If a Provider supports the *plan resources*, a CAMP Consumer can create a new *plan resource* without creating an *assembly resource* by supplying the contents of, or a reference to, either a PDP or a Plan file to the *plan resource* collection in an HTTP POST request (see Section 7.2, “[Registering a Plan](#)”). The loaded *plan resource* model might then appear as follows:

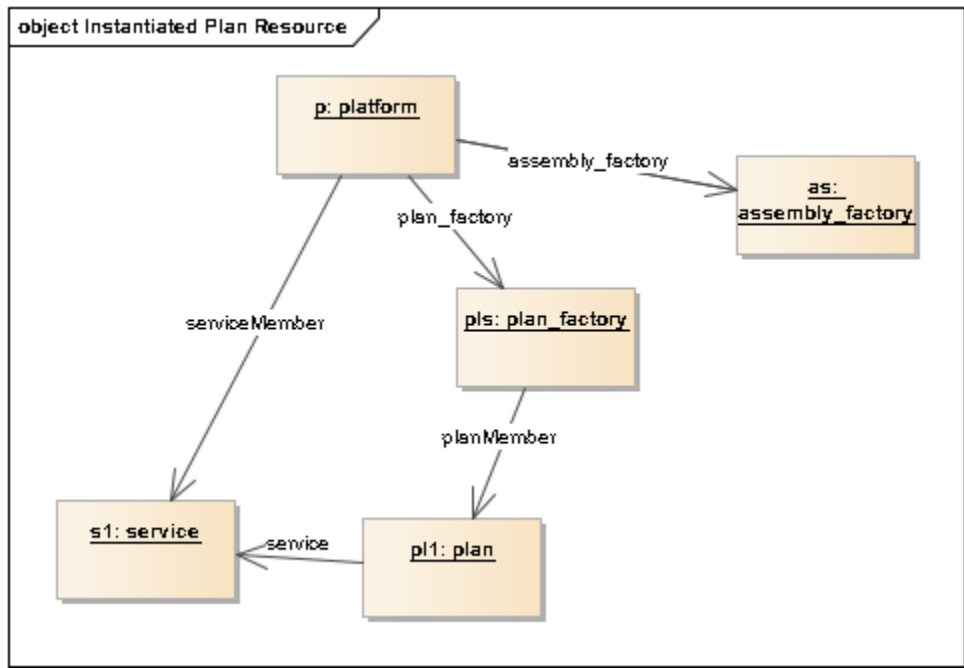


Figure 3-3: Loaded Plan Resource

A CAMP Consumer can create a new *assembly resource* from an existing *plan resource* by providing the reference to that *plan resource* to the *assembly_factory resource* in a HTTP POST request (see Section 7.1.1, “[Deploying an Application by Reference](#)”).

Using this two-step process, the loaded *assembly resource* model would appear the same as when using the one-step process, as shown in Figure 3-2.

3.4 Managing an Application Assembly




Figure 3-4: Instantiated Resources

To manage the operation of the application, the Application Administrator interacts with the *assembly resource* and the related *component resources*.

The traversal of the resources in the model can be accomplished by following the navigation arrows on the associations in these object instance diagrams, from each resource to the other resources it depends on.

The Application Administrator can observe real-time operational metrics through *sensor resources* on *assembly resources* and *component resources*. In response to these metrics, the Application Administrator — or an automated process such as a management system — can affect changes to those resources through the *operation resources* linked from those same resources.

3.5 Removing Assemblies

When finished working with an application, an Application Administrator can delete an *assembly resource* using a DELETE request. The CAMP platform will typically soon thereafter remove the *assembly resource* and all associated resources which are dedicated to that *assembly resource*. Where such a resource is not removed immediately, for example, when it is in the process of shutting down, it ought to present a representation skew of DESTROYING in the interim.

When the original *plan resource* is no longer needed, an Application Administrator can delete it using a DELETE request. Again, the CAMP platform will typically delete the *plan resource* and all associated resources which are dedicated to that *plan resource*. Where this deletion is accepted but not immediate, such as because an *assembly resource* is in use that references the *plan resource*, again the CAMP platform ought to present a representation skew of DESTROYING for the resources being deleted.

4 Platform Deployment Package

The Platform Deployment Package (PDP) and the Plan file (see Section 4.2, “[Plan Overview](#)”) ensure portability across platforms. It can be exported by an instance of the platform to import into another instance of the platform. It can also be created by an Application Development Environment and/or other tools.

4.1 PDP Package Structure

A PDP is an archive which contains a Plan file named `camp.yaml` at the root of the archive. A PDP archive MAY include other files related to the application including, but not limited to, language-specific bundles, resource files, application content files such as web archives, database schemas, scripts, source code, localization bundles, and icons; and metadata files such as manifests, checksums, signatures, and certificates. [\[PDP-01\]](#)

4.1.1 Supported Archive Formats

A Provider SHALL support the following archive file formats for a PDP:

- A PDP as a ZIP archive [\[ZIP\]](#) [\[PDP-02\]](#)
- A PDP as a TAR archive [\[TAR\]](#) [\[PDP-03\]](#)
- A PDP as a GZIP [\[RFC1952\]](#) compressed TAR archive [\[PDP-04\]](#)

Providers MAY support additional archive file formats for the PDP. [\[PDP-05\]](#)

4.1.2 Validating Integrity

A PDP MAY contain a manifest file, named `camp.mf`, at the root of the archive. [\[PDP-06\]](#) This file contains SHA256 [\[SHA256\]](#) digests of some or all files in the package. A Provider SHOULD reject a PDP if any digest listed in the manifest does not match the computed digest for that file in the package. [\[PDP-07\]](#)

A PDP MAY contain a certificate, named `camp.cert`, at the root of the archive. [\[PDP-08\]](#) This file contains a signed SHA256 digest for the manifest file and the corresponding X.509 certificate. A Provider SHOULD reject any PDP for which the signature verification fails. [\[PDP-09\]](#)

The format of the manifest file and the certificate file SHALL be as defined by the OVF specification [\[OVF\]](#). [\[PDP-10\]](#)

4.2 Plan Overview

The Plan provides a description of the artifacts that make up an application, the services that are required to execute or utilize those artifacts, and the relationship of the artifacts to those services. Plans can be represented in two ways, either as YAML files or as CAMP resources. The examples in this section show Plans as YAML files.

Example 1: Minimal Plan describing an application consisting of a single RPM file

```
00 camp_version: CAMP 1.2
01 artifacts:
02   -
03     type: org.rpm:RPM
04     content: { href: my-app.rpm }
```

The above example describes an application that consists of a single RPM ([RPM Package Manager](#)) package, named “my-app.rpm”, which exists at the root of the PDP archive.

4.2.1 Types

Plans can contain descriptions of artifacts, services and their relationships. However, it is outside the scope of this specification to provide detailed definitions of these entities. Instead Plans use 'type' nodes to identify these things. 'Type' nodes are Strings that describe entities that are managed by CAMP, but whose value and semantics are defined outside the CAMP specification. For example, a group of PaaS providers could agree to use the artifact type "org.rpm:RPM" to identify RPM packages. Line 03 in Example 1, above, is an example of the use of such a type.

To promote portability, both providers and consumers of the CAMP API are encouraged to namespace-qualify the types that they use. For example, if a PaaS provider supports a requirement type that expresses the relationship "deploy on a Spring container", the value "com.paas-r-us.spring.DeployOn" is preferable to the value "DeployOn", as the latter is likely to collide with similar types.

In addition to defining the labels for artifacts, services, and their relationships it is expected that those individuals and organizations that define such labels will also define additional attributes that qualify and constrain the entity that is referenced.

Note that this specification does not specify a normative mechanism for Providers to advertise their supported type nodes.

4.2.2 Requirement Specifications

Although Example 1 is a complete, CAMP-conformant Plan, it is somewhat abstract. It essentially says "this application is made up of the following RPM file". It does **not** say anything about what a CAMP Provider is supposed to do with the RPM file. When presented with the Plan in Example 1, a Provider is free to do whatever it likes with the artifact. The obvious action is to install the RPM file (the fact that there is an obvious action is what makes Example 1 workable) but not all artifacts will necessarily have such obvious actions.

Requirement Specifications allow Application Developers to specify what the Provider should do with an artifact.

Example 2: Expanded Plan describing details of how to install the RPM

```
00 camp_version: CAMP 1.2
01 artifacts:
02   -
03     type: org.rpm:RPM
04     content: { my-app.rpm }
05     requirements:
06       -
07         type: org.rpm:Install
08         org.rpm.installopts.excludedocs: true
```

Example 2 adds a Requirement Specification (lines 07-08) to indicate (through the type value of "org.rpm:Install" on line 07) that target Providers are to install the RPM. Furthermore it indicates, on line 08, something about how the Provider should install the RPM (i.e. in a way that excludes documentation).

The correct processing of this requirement is predicated upon the Providers understanding of the structure and semantics of the "org.rpm:Install" type. The Provider has to know that this type indicates that the parent RPM artifact is to be installed on a Linux instance and that Requirement Specifications of this type might contain an org.rpm.installopts.excludedoc node whose value is a boolean that indicates whether files marked as documentation should be installed. It is assumed that the semantics associated with "org.rpm:Install" requirement type are documented, and that this documentation also describes the value space and semantics of the org.rpm.installopts.excludedocs node.

4.2.3 Service Specifications

Example 2 is more specific than Example 1, but it is still silent about what kind of Linux instance we want the RPM installed on. When presented with the Plan in Example 2, a Provider is free to install the RPM on any kind of Linux instance or even, hypothetically, a non-Linux operating system that supports RPM.

Service Specifications allow Application Developers to constrain or outline the services that can be used to support the requirements they have specified.

Example 3: Expanded Plan that includes a Service Specification

```
00 camp_version: CAMP 1.2
01 artifacts:
02   -
03     type: org.rpm:RPM
04     content: { my-app.rpm }
05     requirements:
06       -
07         type: org.rpm:Install
08         org.rpm.installopts.excludedocs: true
09         fulfillment:
10           characteristics:
11             -
12               type: com.example:Linux
13               com.example.linux.kernelVersion: [3.9.6]
14               org.iaas.bitsize: 64
```

Example 3 adds a Service Specification to the Requirement Specification from Example 2. This Service Specification indicates that, wherever the Provider decides to install the RPM, it has to be a 64-bit Linux instance running a kernel of version 3.9.6.

As in the case of our Requirement Specification, the correct processing of this Plan depends upon the Providers ability to understand Characteristic Specifications of the type “com.example:Linux”. The Provider needs to know that characteristics of this type can contain, among other possible nodes, the com.example.linux.kernelVersion and org.iaas.bitsize nodes as well as understand the allowed values and the semantics of these nodes.

Although Example 3 is fairly specific (“this application is made up of an RPM file which is to be installed, excluding documentation, on a 64-bit Linux instance running a kernel of version 3.9.6”), it is still somewhat abstract in that it does not reference the specific Linux instance that the RPM is to be installed on. If an Application Developer wishes to explicitly identify the exact Linux instance on which their RPM is to be installed, they can do so using the href node of a Service Specification.

Note: See [Appendix C](#) for suggested version range values to use when no prevailing scheme already exists for the type.

Example 4: Concrete Plan with service resource reference

```
00 camp_version: CAMP 1.2
01 artifacts:
02   -
03     type: org.rpm:RPM
04     content: { my-app.rpm }
05     requirements:
06       -
07         type: org.rpm:Install
08         org.rpm.installopts.excludedocs: true
09         fulfillment:
10           href: http://example.org/my_paas/services/8675309
```

Example 4 amends the Service Specification introduced in Example 3 to reference an instance of a *service resource* provided by the CAMP implementation at “example.org”. Since the Application Developer is calling for the use of a specific service, it is no longer necessary to indicate to the Provider the general characteristics of the services that are suitable for fulfilling the “org.rpm:Install” requirement. Therefore, the Characteristic Specification in Example 3 has been removed

Note that the specificity in Example 4 comes at the expense of portability. Due to the reference to a specific resource, it is doubtful that the Plan in Example 4 could be successfully deployed on any CAMP instance other than the one at “example.org” whereas the Plan in Example 3 can be deployed on any CAMP instance that supports the “org.rpm:RPM” artifact type, the “org.rpm:Install” requirement type, and the “com.example:Linux” characteristic type. This tradeoff between specificity and portability is a design

feature of CAMP Plans. There are cases (e.g. when developing an application) in which it makes sense to target a Plan for a specific CAMP instance and there are cases (e.g. when moving an application to a different provider) when it makes sense to make the Plan as generic as possible. Plans are designed to allow the Application Developer to be as specific or generic as necessary to accomplish their particular task.

4.2.3.1 Shared Services

There are situations in which an application can have two or more artifacts that need to share the same runtime instance of a service.

Example 5: Plan with shared Service Specification

```
00 camp_version: CAMP 1.2
01 artifacts:
02   -
03     type: com.java:WAR
04     content: { href: vitaminder.war }
05     requirements:
06       -
07         type: com.java:HostOn
08         com.java.servlet.contextName: "/vitaM"
09         fulfillment:
10           ...
11       -
12         type: com.java.jdbc:ConnectTo
13         fulfillment: id:db
14     -
15     type: org.sql:SqlScript
16     content: { href: vitaminder.sql }
17     requirements:
18       -
19         type: org.sql:ExecuteAt
20         fulfillment: id:db
21 services:
22   -
23     id: db
24     characteristics:
25       -
26         type: org.storage.db:RDBM
27       ...
28     -
29     type: org.storage.db:Replication
30     ...
31     -
32     type: org.iso.sql:SQL
```

The above example describes an application with two components, a WAR file and an SQL script. In the case of this particular application, the SQL script is used to initialize the database that will be used by the WAR file. The components created from the two artifacts need to share a common database instance or the application will not work. Lines 23-33 describe the shared target database service. Line 23 is an 'id' node with the value 'db'. This node is used as the target for the 'fulfillment' nodes on lines 13 and 20. The common use of the "id:db" value in lines 13 and 20 indicates that, whatever service used to satisfy the Service Specification in lines 23-33, it will be shared by the components that are created by resolving the requirements on lines 12-13 and lines 19-20.

4.2.3.2 Service Frameworks

There are situations in which the artifacts of an application are dynamically added (e.g. via a git [Git] push operation) after the creation of a "service framework" on which these artifacts can be deployed. Such a framework can be specified via a Plan that contains Service Specifications but no Artifact Specifications.

Example 6: Plan with only Services Specifications

```
camp_version: CAMP 1.2
services:
-
  name: Rails Runtime
  characteristics:
  -
    type: org.ruby-lang:Ruby
    ...
  -
    type: org.rubyonrails:Rails
    ...
-
  name: Database
  characteristics:
  -
    type: org.storage.db:RDBM
    ...
-
  name: Git Repo
  characteristics:
  -
    type: com.git-scm:GIT
    ...
```

The above example specifies a set of services onto which the user can deploy Rails components by pushing them to the git repository that will be created as a result of deploying this Plan.

4.2.4 Names, Description, and Tags

Plans, artifacts and services can be decorated with names, descriptions, and tags. CAMP implementations can use this information when creating the resources that correspond to these entities. For example, the following Plan file:

Example 7: Plan with names, descriptions, and tags

```
name: Mike's Drupal Instance
description: Drupal 6.28
tags: [ PHP, Drupal6, mikez ]
camp_version: CAMP 1.2
artifacts:
-
  type: net.php:Module
  content:
    href: ftp://ftp.drupal.org/files/projects/drupal-6.28.tar.gz
...
```

when successfully registered, could result in the creation of the following *plan resource*:

```
{
  "type": "plan",
  "uri": "http://uswest.paas-r-us.com/camp/plan/101",
  "name": "Mike's Drupal Instance",
  "description": "Drupal 6.28",
  "tags": [ "PHP", "Drupal6", "mikez" ],
  ...
}
```

4.3 Plan Schema

A Platform Deployment Package (PDP) SHALL contain a single Plan file. [PDP-11] The Plan file SHALL be located at the root of the PDP archive. [PLAN-01] The Plan file SHALL be named "camp.yaml".

[PLAN-02] The Plan file SHALL conform to YAML 1.1 [YAML 1.1]. [PLAN-08] The Plan file SHALL conform to the description provided in this section. [PLAN-09]

4.3.1 General Nodes

Plans, Artifact Specifications, and Service Specifications MAY contain the following nodes: [PDP-30]

4.3.1.1 name

Type: String

Required: false

This node expresses the human-readable name of the Plan or Specification. Providers MAY reflect the value of this attribute in the names of any resources that are created in the processing the Plan. [PDP-14]

4.3.1.2 description

Type: String

Required: false

This node expresses the human-readable description of the Plan or Specification. Providers MAY reflect the value of this attribute in the descriptions of the resources that are in the processing the Plan. [PDP-15]

4.3.1.3 tags

Type: String[]

Required: false

This node expresses an array of human-readable tags for the Plan or Specification. Providers MAY reflect the values of this attribute in the tags of the resources that are created in the processing of the Plan. [PDP-16]

4.3.2 Plan

A Plan defines the structure the elements in a Plan file or resource. A Plan file SHALL contain a single instance of a Plan. [PLAN-03] A Plan has the following pseudo schema:

```
name: String ?
description: String ?
tags: String[] ?
camp_version: String
origin: String ?
artifacts: ArtifactSpecification[] ?
services: ServiceSpecification[] ?
```

In addition to the general nodes, a Plan contains the following nodes:

4.3.2.1 camp_version

Type: String

Required: true

The value of this node expresses the version of the CAMP specification to which the Plan conforms. For Plans that conform to this document, the value of this node SHALL be as defined in Section 1.8 "Specification Version". [PLAN-05]

4.3.2.2 origin

Type: String

Required: false

The value of this node specifies the origin of the Plan. For example, when exporting a *plan resource* into a PDP, a Provider might use the URL of its *platform resource* for this value. Alternatively, an Application Development Environment could use its name and version.

4.3.2.3 artifacts

Type: ArtifactSpecification[]

Required: false

This node lists the artifacts that comprise the application described by the Plan. For portability reasons, Providers are cautioned against regarding the order of the elements in this array as significant.

4.3.2.4 services

Type: ServiceSpecification[]

Required: false

This node describes the services that the application described by the Plan requires in order to function. For portability reasons, Providers are cautioned against regarding the order of the elements in this array as significant.

4.3.3 ArtifactSpecification

An ArtifactSpecification describes an artifact of the application. The artifact MAY be contained within the PDP or MAY exist in some other location. [PDP-22]

An ArtifactSpecification has the following pseudo schema:

```
name: String ?
description: String ?
tags: String[] ?
type: String
content: ContentSpecification
requirements: RequirementSpecification[] ?
```

In addition to the general nodes, an ArtifactSpecification contains the following nodes:

4.3.3.1 type

Type: String

Required: true

The value of a type node specifies the type of an artifact.

Note: Values for a type node are not defined by this specification. See Section 4.2.1, “Types”.

4.3.3.2 content

Type: ContentSpecification

Required: true

This node identifies the location of the content of the artifact described by this Artifact Specification. See Section 4.3.4, “ContentSpecification”, for details.

4.3.3.3 requirements

Type: RequirementSpecification[]

Required: false

This array specifies the ways in which the artifact described by this Artifact Specification engages with the services provided by the platform. See Section 4.3.5, “RequirementSpecification”, for details. For portability reasons, Providers are cautioned against regarding the order of the elements in this array as significant.

4.3.4 ContentSpecification

A ContentSpecification defines the content of a component. A ContentSpecification has one of two, mutually exclusive, nodes: href or data. It has the following pseudo schema:

```
href: URI
```

or

```
data: String
```

When href is used in a ContentSpecification its value is interpreted as follows:

- Providers SHALL support the “https” URI scheme as defined in RFC 2818 [RFC2818]. [PDP-27] A Provider MAY support additional URI schemes listed at <http://www.iana.org/assignments/uri-schemes/uri-schemes.xhtml>. [PDP-28]
- URL’s with the special scheme “pdp:” are interpreted as files contained in the PDP.
 - If the path segment (after the “pdp:”) begins with a “/” it is an absolute path.
 - If the path segment is “!” (i.e. the URL is “pdp:!”), the reference is to the PDP archive itself. This is useful in making an existing deployment package (such as a WAR) function as a PDP.
 - For any other path segment, the path is relative to the location of the file which contains the Content Specification, subject to the guidelines below.
 - Where the path segment contains the special character “!”, it is treated as a delimiter to look for the path to the right of “!” inside the archive at the path to the left of the “!”. Providers SHALL understand this delimiter and SHALL NOT resolve any content if the archive format is unsupported. [PDP-29] Consumers SHALL follow the syntax and semantics described here when using URIs with a “pdp” scheme. [PLAN-07] For example “pdp:/certs.zip!/id_rsa.pub” refers to a file “id_rsa.pub” contained at the root of a “certs.zip” file located at the root of the PDP, and is valid only on platforms which support the ZIP format in conjunction with “!”. On other platforms the link will not be resolved.
- Where the value is not a URI, it is interpreted as a “pdp:” protocol link, as though it were preceded by “pdp:/”.

Example 8: A Plan describing an application consisting of the contents of the PDP

```
00 artifacts:
01   -
02     type: org.oasis-open.tosca:CSAR
03     content: { href: pdp:! }
04     requirements:
05       -
06         type: com.oasis-open.tosca:DeployOn
```

The above example illustrates the use of the “pdp:!” construct wherein the content being referenced (on line 03) is the PDP itself. In this case the PDP is also an OASIS TOSCA v1 Cloud Service Archive.

4.3.5 RequirementSpecification

A RequirementSpecification describes the relationship between an artifact and a platform. Providers are expected to use the information in a RequirementSpecification to determine what actions to perform on/with the artifact described in the containing Artifact Specification.

A RequirementSpecification has the following pseudo schema:

```
type: String
fulfillment: (String | ServiceSpecification) ?
```

4.3.5.1 type

Type: String

Required: true

The value of this node defines the relationship of the artifact that contains this RequirementSpecification to a service. For example, “com.java:HostOn”. See Section 4.2.1, “Types”, for a general description of the definition and treatment of these values.

It is expected that RequirementSpecifications will contain extension nodes that modify or provide additional information about the relationship that they describe. The value space and semantics of these extensions ought to be part of the definition of the value used in the “type” node. For example, the definition of the “com.java:HostOn” relationship might define a “com.java:contextPath” node whose value specifies the desired context path for the artifact when it is deployed on its selected service.

4.3.5.2 fulfillment

Type: String or ServiceSpecification

Required: false

The value of this node either describes, or references a description of, the other party in the relationship (i.e. the service) defined by this RequirementSpecification. In the case where this node references a description, the value is a String that corresponds to the id node of a ServiceSpecification (e.g. “id:db”). In the case where this node contains the description, the value is a ServiceSpecification. See Section 4.3.6, “ServiceSpecification”, for details.

4.3.6 ServiceSpecification

A ServiceSpecification describes a service used by the application. This description is not intended to be a complete list of every detail of the service but, instead, an enumeration of those facets that, for whatever reason, are important to the application described by the Plan. Providers are expected to use the information in a ServiceSpecification to select an appropriate service for resolving the containing RequirementSpecification.

A ServiceSpecification has the following pseudo schema:

```
name: String ?
description: String ?
tags: String[] ?
id: String ?
href: URI ?
characteristics: CharacteristicSpecification[]
```

In addition to the general nodes, a ServiceSpecification contains the following nodes:

4.3.6.1 id

Type: String

Required: false

The value of this node serves as an anchor for intra-Plan references. See Section 4.3.5.2, “fulfillment”, for information on how this anchor is used. Plans SHALL use id values that are unique within the scope of the Plan. [PLAN-06]

4.3.6.2 href

Type: URI

Required: false

The value of this node is a reference to a *service resource* (see Section 5.13, “service Resource”) that resolves the service described by this ServiceSpecification. If a Consumer includes this node in a Plan, the value of this node SHALL reference a Consumer-visible resource within the target Platform. [RMR-01]

4.3.6.3 characteristics

Type: CharacteristicSpecification[]

Required: true

This array provides the characteristics of the service described by this ServiceSpecification. See Section 4.3.6, “[CharacteristicSpecification](#)”, for details. For portability reasons, Providers are cautioned against regarding the order of the elements in this array as significant.

4.3.7 CharacteristicSpecification

A CharacteristicSpecification describes a desired characteristic or capability of a service. It has the following pseudo schema.

```
type: String
<name>: <value> *
```

A CharacteristicSpecification is a node with a single, CAMP-defined key named type. A CharacteristicSpecification can have zero or more additional nodes whose values can be any valid YAML value.

The inclusion of a CharacteristicSpecification in a ServiceSpecification indicates that the characteristics being described are significant to the application, but the degree of this significance (e.g. “absolutely necessary” versus “would be nice to have”) is not indicated.

4.3.7.1 type

Type: String

Required: true

The value of this node defines the characteristic being described by this CharacteristicSpecification. For example, “com.java:ServletContainer”. See Section 4.2.1, “[Types](#)”, for a general description of the definition and treatment of these values.

It is expected that CharacteristicSpecifications will contain extension nodes that modify or provide additional information about the characteristic that they describe. The value space and semantics of these extensions ought to be part of the definition of the value used in the type node. For example, the definition of the “org.rubyonrails:Rails” characteristic might define a “org.rubyonrails:version” node whose value specifies the version of Rails provided by the service.

Note: See [Appendix C](#) for suggested version range values to use when no prevailing scheme already exists for the type.

5 Resources

The following sub-sections describe the resources defined by this specification.

When supporting such a Resource, a Provider SHALL implement it and serialize it as described in the corresponding sub-section. [RE-70]

A Consumer SHALL serialize Resource data in its requests based on the definition of this Resource as described in the corresponding sub-section. [RE-71]

5.1 Required Attributes

If the Required boolean constraint for an attribute of a resource type has a value of "true", then a resource of this type SHALL have the attribute present. [RE-06] If the value is "false" then the resource is valid with or without the attribute present.

5.2 Attribute Types

Resource attributes are defined using the following types:

5.2.1 Boolean

As defined by JSON [RFC4627], a token having a literal value of either `true` or `false`. The use of this type is indicated in metadata by an *attribute_definition resource* with an *attribute_type* value of "Boolean".

5.2.2 String

A UNICODE string as defined by JSON [RFC4627]. The use of this type is indicated in metadata by an *attribute_definition resource* with an *attribute_type* value of "String".

5.2.3 Number

As defined by Section 2.4 of JSON [RFC4627]. The use of this type is indicated in metadata by an *attribute_definition resource* with an *attribute_type* value of "Number".

5.2.4 URI

A String (see above) that conforms to the syntax defined in RFC 3986 [RFC3986]. The use of this type is indicated in metadata by an *attribute_definition resource* with an *attribute_type* value of "URI".

5.2.5 Timestamp

A String (see above) that conforms to the syntax defined in ISO 8601 [ISO 8601:2004]. Consumers and Providers SHALL express Timestamps in UTC (Coordinated Universal Time), with the special UTC designator ("Z"). [RE-65] The use of this type is indicated in metadata by an *attribute_definition resource* with an *attribute_type* value of "Timestamp".

5.3 CAMP Resource Type Inheritance

Each CAMP resource has a resource type associated with it, indicated by the *type_definition* URI in the resource's metadata section (see Section 5.4.7.1, "*type_definition*"). A resource's type defines the attributes of that resource along with the constraints and semantics of those attributes. Resource types form an inheritance hierarchy with *camp_resource* (See Section 5.4, "*camp_resource Resource*") at its root. When a resource type (sub-type) inherits from another resource type (super-type), the sub-type inherits, and therefore includes, all the super-type's attributes along with its constraints and semantics. A

sub-type can add additional attributes not present in its super-type(s). A sub-type MAY further restrict the constraints of an attribute inherited from its super-type(s). [MO-01] A sub-type SHALL NOT loosen the constraints of an attribute inherited from its super-type(s). [MO-02] As a consequence, a resource of a super-type can always be substituted with a resource of any of its sub-types. A resource type MAY inherit from more than one super-type. [MO-03] If there is an attribute name collision when a sub-type inherits from multiple super-types, the inherited attributes of the same name SHALL NOT contradict the constraints and semantics of the attributes defined in its super-types. [MO-04]

5.4 camp_resource Resource

All CAMP resources SHALL inherit directly or indirectly from this resource. [MO-05] This resource contains the following attributes:

5.4.1 uri

Type: URI

Required: true

This attribute expresses the URI of the resource.

5.4.2 name

Type: String

Required: true

This attribute expresses the human-readable name of the resource.

5.4.3 description

Type: String

Required: false

This attribute expresses the human-readable description of the resource.

5.4.4 tags

Type: String[]

Required: false

This attribute is an array of String values that may be assigned by the provider or the user. These values can be used for keywording and terms-of-interest.

5.4.5 representation_skew

Type: String

Required: false

The `representation_skew` attribute expresses the relationship between the information presented in the resource and the status of the platform implementation artifacts that are represented by that resource (see Section 2.7, “[Representation Skew](#)”). It is an optional, enumerated String. If present, `representation_skew` SHALL have one of the following values: [RE-11]

- “CREATING” – describes a resource that is in the process of being created. The client can expect that the resource will have a skew of “NONE” once this process has completed.
- “NONE” – is an assertion by the CAMP server that the information in the resource is an accurate representation of the underlying platform implementation. Absent some action by the client or some other event (e.g. platform shutdown), a resource with a skew of NONE can be expected to remain in synch with the platform implementation.

- “UNKNOWN” – indicates that the CAMP server cannot accurately depict the aspect of the platform implementation represented by this resource. Users can attempt to address the underlying issues(s) by manipulating this and/or other resources as specified by the API.
- “DESTROYING” – describes a resource that is in the process of being destroyed. The client can expect that the resource will cease to exist once this process has completed.

The absence of the `representation_skew` attribute is semantically equivalent to a value of “NONE”.

The value of the `representation_skew` attribute applies only to the resource of which it is part. The skew of any resources that are contained by another resource (e.g. in the manner in which the *assembly resource* contains *component resources*) is conveyed by the individual `representation_skew` of those sub-resources and not aggregated or “rolled up” into the containing resource.

The value of the `representation_skew` attribute affects the availability of the HTTP methods for that resource. For example, resources with a `representation_skew` value of CREATING might support the GET, HEAD and DELETE methods, but no other HTTP methods. The following table lists the methods that SHALL be supported for each `representation_skew` value. [RE-12]

<code>representation_skew</code> value	Methods Available
CREATING	GET, DELETE
NONE	All supported methods for that resource.
UNKNOWN	All supported methods for that resource.
DESTROYING	GET

Table 5-1: `representation_skew` Available Methods

For each `representation_skew` value, CAMP Providers MAY support HTTP methods in addition to those listed in the corresponding row of Table 5-1 [RE-13]

5.4.6 external_management_resource

Type: URI

Required: false

A URI to an external management interface to manage the entity represented by this resource (see Section 2.8 “External Management Resources”). The entity referred to by this attribute is platform dependent and requires external documentation to understand its meaning.

5.4.7 metadata

Type: object

Required: true

The value of this attribute specifies the metadata for the resource. Three attributes are defined for metadata: `type_definition`, `mutable`, and `consumer_mutable`. Metadata for a resource is extensible; a Platform Provider MAY provide additional metadata attributes. [RE-81]

5.4.7.1 type_definition

Type: URI

Required: true

The value of this attribute is a reference to the *type_definition resource* that describes this resource.

5.4.7.2 mutable

Type: array

Required: false

The value of this attribute is a list of the attributes of the resource that are mutable. If a particular attribute is not listed, it indicates that the value(s) of the attribute, once set, SHALL NOT change for the lifetime of the resource. [RE-07] If a particular attribute is listed, it indicates that the value of the attribute MAY change due to the actions or activity of either the Provider or the Consumer. [RE-08]

Elements of this array use JSON Pointer [RFC6901] syntax to identify the attribute that is mutable.

5.4.7.3 consumer_mutable

Type: array

Required: false

The value of this attribute is a list of attributes of the resource that are consumer-mutable. If a particular attribute is not listed, it indicates that the value(s) of the attribute SHALL NOT be changed by Consumers. [RE-09] If a particular attribute is listed, it indicates that Consumers MAY change the value of the attribute. [RE-10] Note that absence of an attribute in the list does not preclude the Provider from changing the value of the attribute.

The values in this list are only relevant for mutable attributes. A Provider SHALL NOT list an attribute in the consumer_mutable array if it is not present in the mutable array. [RE-82]

Elements of this array use JSON Pointer [RFC6901] syntax to identify the attribute that is consumer-mutable.

5.5 HTTP Method Support

As described in Section 6.1, “Transfer Protocol”, Consumers use HTTP [RFC2616] to interact with CAMP-defined resources. To foster interoperability it is necessary to define the HTTP methods supported by each resource. Note that a requirement on the Provider to support a particular HTTP method on a resource does not ensure that all requests to that resource using that method will succeed; it simply guarantees that the Provider will not fail such requests with a 405 (Method Not Allowed) error.

Providers SHALL support the HTTP GET method on all of the resources defined in this section. [RE-53] Providers SHALL support the HTTP PUT and PATCH methods on resources that have at least one consumer-mutable attribute. [RE-83] Note that Providers are required to provide metadata on the consumer mutability of all attributes. Requirements for the support of additional HTTP methods are outlined in the descriptions of each resource below. Providers MAY elect to support additional HTTP methods in addition to those described here. [RE-54]

There are interactions between various HTTP methods and the mutable and consumer-mutable attributes that warrant some explanation. Note that it is possible for a Consumer to cause the Provider to change the values of some mutable attributes in response to a POST operation. For example, a successful POST to an *operation resource* could cause the Provider to change the value of the status attribute of the target resource for that operation. However, this does **not** mean that the status attribute of that target resource is consumer-mutable. Consumer mutability refers to the ability of the Consumer to **directly** alter the value of one or more attributes via the PUT or PATCH methods (see Section 7.4, “Updating Resources”).

5.6 collection Resource

A *collection resource* is a representation of an improper subset of the resources in a collection. *collection resources* have the following pseudo schema:

```

{
  "uri": URI,
  "name": String,
  "description": String ?,
  "tags": String[] ?,
  "representation_skew": String ?,
  "external_management_resource": URI ?,
  "collection_type": URI,
  "total_items": Number,
  "items_per_page": Number,
  "start_index": Number,
  "items": camp_resource[],
  "metadata": object
}

```

If a Provider chooses to return the ETag header in its responses, the value of this header SHALL refer to the entire collection. [RE-84] Sort order and page size does not affect the ETAG value.

collection resources have the following attributes in addition to those defined in *camp_resource*:

5.6.1 collection_type

Type: URI

Required: true

The value of this attribute is a reference to the *type_definition resource* that describes the type of the resources in this collection. All of the resources in this collection will either be of the type described by this *type_definition resource* or of a type that inherits, directly or indirectly, from that type. A Provider MUST only return representations of *collection resources* whose *items* array contains resources whose type is either that described by the *type_definition resource* referenced by the *collection_type* attribute or a type derived, directly or indirectly, from that type. [RE-80]

5.6.2 total_items

Type: Number

Required: true

The value of this attribute is the total number of resources in the collection, as seen by the Provider. This may be greater than the number of resources returned in the *collection resource*, for example if the GET request specified either of the *start_index* or *max_page* query string parameters to request pagination. See Section 7.3.4, “[Pagination of a Collection](#)”, for more details. The value of *total_items* attribute SHALL be a non-negative integer. [RE-85]

5.6.3 items_per_page

Type: Number

Required: true

The value of this attribute is the number of resources returned in the *collection resource*. This may be fewer than the total number of resources in the collection, for example if the GET request specified either of the *start_index* or *max_page* query string parameters to request pagination. See Section 7.3.4, “[Pagination of a Collection](#)”, for more details. The value of the *items_per_page* attribute SHALL be a non-negative integer. [RE-86]

5.6.4 start_index

Type: Number

Required: true

The value of this attribute is the zero-based, ordinal number of the first element returned from the collection, after sorting. If the *start_index* query parameter was used on the HTTP GET request, the

value of this attribute SHALL equal the value of that query parameter. [RE-87] See Section 7.3.4, “[Pagination of a Collection](#)”, for more details. The value of the `start_index` attribute SHALL be a non-negative integer. [RE-88]

5.6.5 items

Type: `camp_resource[]`

Required: true

The value of this attribute is an array of objects that comprise the resources in the collection. The objects in the `items` array are independently accessible resources identified by the item's `uri` attribute.

The size of the array SHALL be equal to the value of the `items_per_page` attribute. [RE-89] If the collection is empty, the value of this attribute SHALL be an array with no elements (represented as “`[]`” in JSON). [RE-90]

5.7 platform_endpoints Collection

A Provider MAY concurrently offer multiple instances of the CAMP API. [RE-15] The primary example of why a provider might do this is to simultaneously support two or more incompatible versions/implementations of the CAMP API, but there are many reasons for a provider to offer multiple instances of the CAMP API.

Concurrent instances are supported through the use of multiple *platform resources*. A collection of *platform_endpoint resources* allows Consumers to discover all the instances of the CAMP API that are currently available.

A Provider MAY expose the *platform_endpoints resource* collection and corresponding *platform_endpoint resources* in a way that allows for version discovery before the client has authenticated. [RE-17]

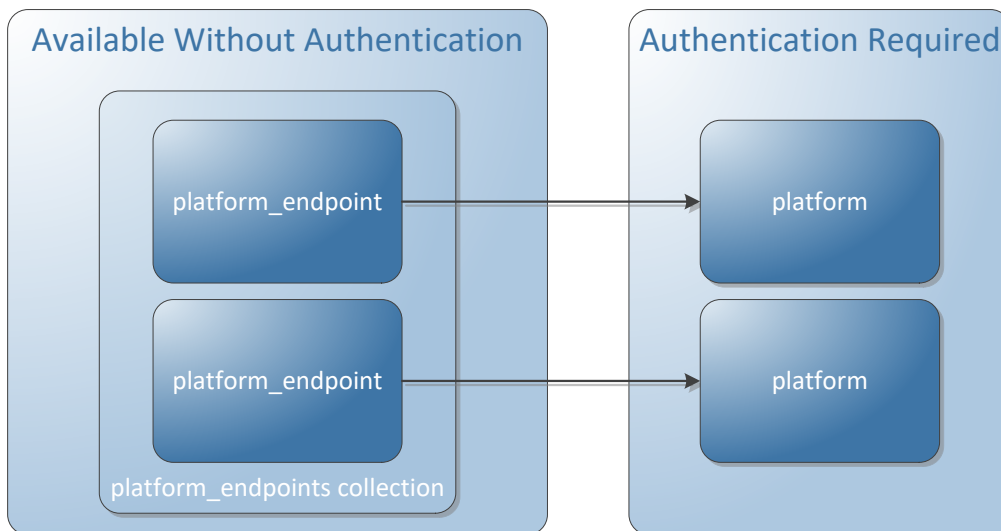


Figure 5-1: Example Implementation

5.8 platform_endpoint Resource

Each *platform_endpoint resource* SHALL refer to exactly one *platform resource*, and indicate the versions supported by the Platform. [RE-20] This specification is deliberately silent about any relationship between resources within different *platform* trees. Each *platform resource* could represent a different CAMP API “view” of the same applications and services. On the other hand, each *platform* could represent a completely independent system.

A *platform_endpoint* resource has the following pseudo schema:

```
{
  "uri": URI,
  "name": String,
  "description": String ?,
  "tags": String[], ?
  "representation_skew": String ?,
  "external_management_resource": URI ?,
  "platform": URI,
  "specification_version": String,
  "backward_compatible_specification_versions": String[] ?,
  "implementation_version": String ?,
  "backward_compatible_implementation_versions": String[] ?,
  "auth_scheme": String ?,
  "metadata": object
}
```

Note: Because of its unique function, future versions of the CAMP specification are cautioned against making non-backwards compatible changes to this resource.

Instances of the *platform_endpoint* resource contain the following attributes:

5.8.1 platform

Type: URI

Required: true

This attribute is the URI of the *platform* resource that this *platform_endpoint* resource describes.

5.8.2 specification_version

Type: String

Required: true

Each *platform* resource is the root of a tree of resources, the syntax and semantics of which conform to one or more versions of the CAMP specification. The value of this attribute is the Specification Version String of the CAMP specification that is supported by the resources rooted in the Platform referenced by the *platform* attribute of this resource.

For Platforms that implement this version of the CAMP specification, the value of this attribute SHALL be as defined in Section 1.8, “Specification Version”. [RE-22]

5.8.3 backward_compatible_specification_versions

Type: String[]

Required: false

The values in this array identify each version of the CAMP specification that is backwards compatible with the current *specification_version* of the Platform (referenced in the *platform* attribute of this resource). The values in this array SHALL be the Specification Version Strings of previous CAMP specification versions. [RE-23]

If this attribute is not present, the version of the CAMP specification implemented by the Platform (referenced in the *platform* attribute of this resource) is not backwards compatible with any previous version of the CAMP specification.

platform_endpoint resources that reference *platform* resources with a *specification_version* value of “CAMP 1.2” SHALL NOT include this attribute because no previous versions are compatible. [RE-24]

5.8.4 implementation_version

Type: String

Required: false

Multiple implementations of the same CAMP specification MAY be offered concurrently. [RE-25] For example, a Provider could offer an initial beta version of “CAMP 1.1” and, later, a production version; allowing a period of overlap for their customers to migrate from the beta to the production version. The value of this attribute is an arbitrary String that expresses the Provider-specific implementation version supported by the resources rooted in the Platform (referenced in the platform attribute of this resource).

5.8.5 backward_compatible_implementation_versions

Type: String[]

Required: false

The values in this array list the provider-specific implementation versions that are backwards compatible with the implementation version of the Platform (referenced in the platform attribute of this resource). The values in this array are arbitrary Strings that correspond to previous implementation_version Strings.

If this attribute is not present, the implementation version offered by the Platform (referenced in the platform attribute of this resource) is not backwards compatible with any previous implementation versions.

5.8.6 auth_scheme

Type: String

Required: false

The value of the auth_scheme attribute indicates the authentication scheme expected by the referenced Platform. For interoperability reasons, Providers are encouraged to offer at least one of the following three (case sensitive) values:

Value	Description
RFC2617	HTTP Basic Authentication [RFC2617]
RFC6749	OAuth2 [RFC6749]
KEYSTONE-2.0	OpenStack Keystone Authentication. [Keystone]
NONE	No authentication required.

Table 5-2 - auth_scheme values

Providers are allowed to extend this list, and provide values of their own. Absence of this attribute means that no authentication scheme is advertised.

Note: Omitting the auth_scheme attribute is discouraged for interoperability reasons.

Note: If Providers wish to offer multiple authentication schemes, they may use multiple platform_endpoint resources each with a different auth_scheme value.

5.9 platform Resource

The platform resource represents the Consumer's initial view of the accessible resources and deployed entities. This resource has the following pseudo schema:

```

{
  "uri": URI,
  "name": String,
  "description": String ?,
  "tags": String[] ?,
  "type_definition": URI,
  "representation_skew": String ?,
  "external_management_resource": URI ?,
  "supported_format_collection": URI,
  "extension_collection": URI,
  "type_definition_collection": URI,
  "platform_endpoints_collection": URI,
  "specification_version": String,
  "implementation_version": String ?,
  "assembly_factory": URI,
  "service_collection": URI,
  "plan_factory": URI ?,
  "metadata": object
}

```

The *platform resource* contains the following attributes:

5.9.1 supported_format_collection

Type: URI

Required: false

This attribute is a URL reference to a collection of *format resources* that identify the Supported Formats for this Platform. See Section 5.16, “[format Resource](#)”, for details.

For every format that the Platform supports, there SHALL be a *format resource* that describes that format. [\[RE-40\]](#) A *format resource* that describes the required JSON format SHALL appear in the collection of *format resources* referenced by the *platform resource*. [\[RE-41\]](#)

5.9.2 extension_collection

Type: URI

Required: true

This attribute is a URL reference to the collection of *extension resources* that identify the Extensions supported by this Platform. See Section 8.2, “[extension Resource](#)”, for details.

For every available extension, the *extension resource* collection referenced by the *platform resource* SHALL contain an *extension resource* that represents the extension. [\[EX-08\]](#)

5.9.3 type_definition_collection

Type: URI

Required: true

This attribute is a reference to the collection of *type_definition resources* that provide information on the resource types supported by this Platform. See Section 5.17, “[type_definition Resource](#)”, for details.

For every resource type that the Platform supports, the *type_definition resource* collection SHALL contain a *type_definition resource* that represents that resource type. [\[RE-44\]](#)

5.9.4 platform_endpoints_collection

Type: URI

Required: true

This attribute is a URL reference to the collection of *platform_endpoint resources* that enumerate the currently available CAMP implementations. See Section 5.7 “[platform_endpoints Collection](#)”, for details.

5.9.5 specification_version

Type: String

Required: true

Each *platform resource* is the root of a tree of resources, the syntax and semantics of which conform to one or more versions of the CAMP specification. The value of this attribute is the Specification Version String of the CAMP specification that is supported by the resources rooted in this Platform.

For Platforms that implement this version of the CAMP specification, the value of this attribute SHALL be as defined in Section 1.8, “Specification Version”. [RE-26]

The value of this attribute SHALL exactly match the value of the `specification_version` attribute of any *platform_endpoint resource* that references this *platform resource*. [RE-27]

5.9.6 implementation_version

Type: String

Required: false

A Provider MAY choose to offer multiple implementations of the same CAMP specification. [RE-28] For example, a Provider could offer an initial beta version of “CAMP 1.1” and, later, a production version; allowing a period of overlap for their customers to migrate from the beta to the production version. The value of this attribute is an arbitrary String that expresses the Provider-specific implementation version supported by the resources rooted in this Platform.

The value of this attribute SHALL exactly match the value of the `implementation_version` attribute of any *platform_endpoint resource* that references this *platform resource*. [RE-29]

5.9.7 assembly_factory

Type: URI

Required: true

This attribute is a URL reference to the *assembly_factory resource*. The *assembly_factory resource* both enumerates the applications deployed on this platform and acts as a factory for instantiating new applications. See Section 5.10, “[assembly_factory Resource](#)”, for details.

5.9.8 service_collection

Type: URI

Required: true

This attribute is a URL reference to the collection of *service resources* that enumerate the services available to the Consumer on this platform. See Section 5.13, “[service Resource](#)”, for details.

5.9.9 plan_factory

Type: URI

Required: false

This attribute is a URL reference to the *plan_factory resource*. The (optional) *plan_factory resource* both enumerates the plans deployed on this platform and acts as a factory for creating new plans. See Section 5.14, “[plan_factory Resource](#)”, for details.

5.10 assembly_factory Resource

This collection resource acts as a factory for the *assembly resources* on this platform. This resource is a sub-type of the *collection resource* and has the following pseudo schema:

```

{
  "uri": URI,
  "name": String,
  "description": String ?,
  "tags": String[] ?,
  "representation_skew": String ?,
  "external_management_resource": URI ?,
  "collection_type": URI,
  "total_items": Number,
  "items_per_page": Number,
  "start_index": Number,
  "items": assembly[],
  "parameter_definition_collection": URI,
  "metadata": object
}

```

In addition to the methods defined in Section 5.5, “[HTTP Method Support](#)”, Providers SHALL support the HTTP POST method on the *assembly_factory resource* as described in Section 7.1, “[Deploying an Application](#)”. [\[RMR-02\]](#)

The *assembly_factory resource* contains the following, additional attribute:

5.10.1 parameter_definition_collection

Type: URI

Required: true

The value of the *parameter_definition_collection* attribute references a collection of *parameter_definition resources* that describe the parameters accepted by this resource on an HTTP POST request. Each of the *parameter_definition resources* in this collection provides metadata for a parameter as described in Section 5.19, “[parameter_definition Resource](#)”. The *assembly_factory resource* accepts the *pdp_uri*, *plan_uri*, *pdp_file*, *plan_file*, *name*, *description*, and *tags* parameters to create a new *assembly resource* upon a POST. The *assembly_factory resource* SHALL reference a collection of *parameter_definition resources* that describe the *pdp_uri*, *plan_uri*, *pdp_file*, *plan_file*, *name*, *description*, and *tags* parameters. [\[RMR-03\]](#)

5.11 assembly Resource

An *assembly resource* represents an instantiated application at runtime. This resource has the following pseudo schema:

```

{
  "uri": URI,
  "name": String,
  "description": String ?,
  "tags": String[] ?,
  "representation_skew": String ?,
  "external_management_resource": URI ?,
  "component_collection": URI,
  "plan": URI ?,
  "operation_collection": URI ?,
  "sensor_collection": URI ?,
  "metadata": object
}

```

In addition to the methods defined in Section 5.5, “[HTTP Method Support](#)”, Providers SHALL support the HTTP DELETE method on the *assembly resource*. [\[RE-61\]](#) On reception of a DELETE request a Provider SHALL remove the *assembly resource* from the system along with any *component resources* referenced by that *assembly resource*. (i.e. the tree of resources that was created when the application was instantiated). [\[RE-73\]](#) On reception of a DELETE request a Provider SHALL remove the *assembly resource* from the *assembly_factory resource*’s *items* array. [\[RE-74\]](#)

An *assembly resource* contains the following attributes:

5.11.1 component_collection

Type: URI

Required: true

The value of the `component_collection` attribute is a reference to the collection of *component resources* that make up this *assembly resource*. An *assembly resource* SHALL have at least one *component resource*. [RE-39]

5.11.2 plan

Type: URI

Required: false

The value of this optional attribute is a URL reference to the *plan resource* for this *assembly resource*. Providers that support Plans SHALL include this attribute in all *assembly resources*. [RMR-04]

5.11.3 operation_collection

Type: URI

Required: false

This attribute references a collection of *operation resources*. The *operation resources* in this collection represent the operations available for this application.

5.11.4 sensor_collection

Type: URI

Required: false

This attribute references a collection of *sensor resources*. The *sensor resources* in this collection represent the sensors available for this application.

5.12 component Resource

A *component* represents a component of a running application. This resource has the following pseudo schema:

```
{
  "uri": URI,
  "name": String,
  "description": String ?,
  "tags": String[] ?,
  "representation_skew": String ?,
  "external_management_resource": URI ?,
  "assembly_collection": URI,
  "artifact": URI ?,
  "service": URI ?,
  "status": String,
  "related_component_collection": URI ?,
  "operation_collection": URI ?,
  "sensor_collection": URI ?,
  "metadata": object
}
```

In addition to the methods defined in Section 5.5, “HTTP Method Support”, Providers SHALL support the HTTP DELETE method on the *component resource*. [RE-62] A successful DELETE request stops the underlying component, removes the *component resource* from the system, and removes it from the component collection of its containing *assembly resource*.

Each *component resource* contains the following attributes:

5.12.1 assembly_collection

Type: URI

Required: true

The value of the `assembly_collection` attribute is a reference to the collection of *assembly resources* for which this *component resource* is a member.

5.12.2 artifact

Type: URI

Required: false

The value of the `artifact` attribute is a URL reference to the artifact on which this *component resource* is based. This artifact is not a CAMP resource, but a representation of the actual artifact (e.g. WAR file, Ruby gem file, etc.)

The `artifact` attribute and the `service` attribute are mutually exclusive.

5.12.3 service

Type: URI

Required: false

The value of the `service` attribute is a URL reference to the *service resource* on which this *component resource* is based.

The `service` attribute and the `artifact` attribute are mutually exclusive.

5.12.4 status

Type: String

Required: true

The value of this attribute indicates the status of the component represented by the *component resource*. This attribute MAY have one of the following values:

- "RUNNING" – indicates that the component is functioning as expected.
- "ERROR" – indicates that the component has encountered some sort of error. [RE-68]

Providers MAY extend this list with additional values. [RE-69]

The value of this attribute can change in response to the invocation of an operation (see Section 5.20, "operation Resource") or as a result of some change in the underlying system.

As with other attributes, the value of this attribute cannot be construed to accurately reflect the status of the underlying component if the `representation_skew` has a value other than "NONE".

5.12.5 related_component_collection

Type: URI

Required: false

This attribute is a reference to a collection of the *component resources* to which this component is related.

5.12.6 operation_collection

Type: URI

Required: false

This attribute references a collection of *operation resources*. The *operation resources* in this collection represent the operations available for this component.

5.12.7 sensor_collection

Type: URI

Required: false

This attribute references a collection of *sensor resources*. The *sensor resources* in this collection represent the sensors available for this component.

5.13 service Resource

A *service resource* represents a particular configuration of a service available for use by one or more applications. This resource has the following pseudo schema:

```
{
  "uri": URI,
  "name": String,
  "description": String ?,
  "tags": String[] ?,
  "representation_skew": String ?,
  "external_management_resource": URI ?,
  "parameter_definition_collection": URI ?
  "characteristics": [
    {
      "type": String,
      String: String *
    } +
  ] ?,
  "metadata": object
}
```

The *service resource* contains the following attributes:

5.13.1 parameter_definition_collection

Type: URI

Required: false

Some *service resources* can act as factories for the independent creation of components. Sending an HTTP POST request to the URL of such a resource will create a new component. Often there is a need to parameterize such requests.

The value of this attribute references a collection of *parameter_definition resources* that describe the parameters accepted by this resource on an HTTP POST request. Each of the *parameter_definition resources* in this collection provides metadata for a parameter as described in Section 5.19, "[parameter_definition Resource](#)".

The collection of *parameter_definition resources* referenced by this attribute SHALL define parameters to allow setting the 'name', 'description', and 'tags' attributes of any new resource created in the course of interacting with this resource. [\[RE-37\]](#)

If this attribute is present in the resource, Providers SHALL support the POST method on that resource in addition to the methods defined in Section 5.5, "[HTTP Method Support](#)". [\[RE-38\]](#)

5.13.2 characteristics

Type: Array of CharacteristicSpecification

Required: false

The optional `characteristics` attribute describes the capabilities of the service described by the *service resource*. The elements of this array have the same schema as the `CharacteristicSpecification` (described in Section 4.3.7, “[CharacteristicSpecification](#)”) of a Plan.

Note that this specification is deliberately silent about the process of matching the `ServiceSpecifications` in a Plan to the services described by *service resources*. Any correspondence between the information in a Plan’s `ServiceSpecification` and the information in the `characteristics` attribute does not necessarily constitute a contract to resolve the containing requirement with that service, though Providers are free to implement and advertise such contracts if they wish.

5.14 plan_factory Resource

This optional collection resource acts as a factory for the *plan resources* on this platform. This resource is a sub-type of the *collection resource* and has the following pseudo schema:

```
{
  "uri": URI,
  "name": String,
  "description": String ?,
  "tags": String[] ?,
  "representation_skew": String ?,
  "external_management_resource": URI ?,
  "collection_type": URI,
  "total_items": Number,
  "items_per_page": Number,
  "start_index": Number,
  "items": plan[],
  "parameter_definition_collection": URI,
  "metadata": object
}
```

In addition to the methods defined in Section 5.5, “[HTTP Method Support](#)”, Providers SHALL support the HTTP POST method on the *plan_factory resource* as described in Section 7.2, “[Registering a Plan](#)”.
[RMR-05]

The *plan_factory resource* contains the following, additional attribute:

5.14.1 parameter_definition_collection

Type: URI

Required: true

The value of the `parameter_definition_collection` attribute references a collection of *parameter_definition resources* that describe the parameters accepted by this resource on an HTTP POST request. Each of the *parameter_definition resources* in this collection provides metadata for a parameter as described in Section 5.19, “[parameter_definition Resource](#)”. The *plan_factory resource* accepts the `pdp_uri`, `plan_uri`, `pdp_file`, `plan_file`, `name`, `description`, and `tags` parameters to create a new *plan resource* upon a POST. The *plan_factory resource* SHALL reference a collection of *parameter_definition resources* that describe the `pdp_uri`, `plan_uri`, `pdp_file`, `plan_file`, `name`, `description`, and `tags` parameters. [RMR-06]

5.15 plan Resource

This optional resource stores the *plan* for an application. As discussed in Section 2.3, “[Deployment](#)”, this information is supplied to the platform as part of the operation of deploying an application. This resource has the following pseudo schema:


```

{
  "uri": URI,
  "name": String,
  "description": String ?,
  "tags": String[] ?,
  "representation_skew": String ?,
  "external_management_resource": URI ?,
  "camp_version": String,
  "origin": String ?,
  "artifacts": [
    {
      "name": String ?,
      "description": String ?,
      "tags": String[] ?,
      "type": String,
      "content": { "href": URI },
      "requirements": [
        {
          "type": String,
          "fulfillment": {
            "name": String ?,
            "description": String ?,
            "tags": String[] ?,
            "id": String ?,
            "href": URI ?,
            "characteristics": [
              {
                "type": String,
                String: String *
              } +
            ] ?
          } ?
        } +
      ] ?,
    } +
  ] ?,
  "services": [
    {
      "name": String ?,
      "description": String ?,
      "tags": String[] ?,
      "id": String ?,
      "href": URI ?,
      "characteristics": [
        {
          "type": String,
          String: String *
        } +
      ] ?,
    } +
  ] ?,
  "metadata": object
}

```

The schema of the *plan resource* returned from a CAMP Provider SHALL conform to the schema for Plans described in Section 4.3, “Plan Schema”, with the following additional requirements: [\[RMR-07\]](#)

- Representations of the *plan resource* SHALL be serialized as JSON, unless another format is negotiated. [\[RMR-08\]](#)
- Any href attributes of ServiceSpecifications SHOULD refer to a Service resource. [\[RMR-09\]](#)
- All href attributes in the *plan resource* SHOULD be set to a consumer accessible URL. If the original Plan file referred to a local file, the URL indicates where the Provider stored the content. [\[RMR-10\]](#)

- The *plan resource* inherits from the *camp_resource* defined in Section 5.4, “[camp_resource Resource](#)”, and therefore inherits all its attributes.

For example, if the consumer-supplied Plan file describes an artifact with an href pointing to a file contained in a PDP, the platform-supplied *plan resource* will point to a copy of that artifact, such as one hosted at the platform or in an object store.

Support for the *plan resource* is uniform across a CAMP implementation. Regardless of whether a Consumer attempts to create an *assembly resource* by POSTing to the *assembly_factory resource* or creates a *plan resource* by POSTing to the *plan_factory resource*, a Provider that supports *plan resources* and the *plan_factory resource* SHALL create a *plan resource* for every deployed application. [\[RMR-11\]](#)

In addition to the methods defined in Section 5.5, “[HTTP Method Support](#)”, Providers SHALL support the HTTP DELETE method on the *plan resource*. [\[RE-77\]](#) On reception of a DELETE request a Provider SHALL remove the *plan resource* from the system. [\[RE-78\]](#) On reception of a DELETE request a Provider SHALL remove the *plan resource* from the *plan resource* collection. [\[RE-79\]](#)

5.15.1 Advertising Support for the Plan Resource

As an aid to interoperability it is helpful if Consumers can easily discover if a particular Provider supports *plan resources* and the *plan resource* collection. Section 8.2, “[extension Resource](#)”, defines a mechanism for advertising extensions to the CAMP specification. This mechanism is used to advertise support for *plan resources* and the *plan resource* collection.

Providers that support *plan resources* and the *plan resource* collection SHALL advertise such support using the following *extension resource*: [\[RMR-12\]](#)

```
{
  "uri": <as appropriate>,
  "name": "CAMP Plans Extension",
  "description": "indicates support for plan resources",
  "version": "CAMP 1.2",
  "documentation": "http://docs.oasis-open.org/camp/camp-spec/v1.2/camp-spec-
v1.2.pdf",
  "metadata": object
}
```

5.16 format Resource

A *format resource* represents exactly one supported data serialization format. This resource has the following pseudo schema:

```
{
  "uri": URI,
  "name": String,
  "description": String ?,
  "tags": String[] ?,
  "representation_skew": String ?,
  "external_management_resource": URI ?,
  "mime_type": String,
  "version": String ?,
  "documentation": URI
}
```

Instances of the Format resource contain the following attributes:

5.16.1 mime_type

Type: String

Required: true

This attribute contains the mime-type to be used by the Platform in HTTP [RFC2616] compliant content negotiation for this Format. For example: "application/json".

5.16.2 version

Type: String

Required: false

This attribute contains the version identifier of the data serialization format used.

5.16.3 documentation

Type: URI

Required: true

The value of the documentation attribute is a URI reference to a document that describes the format identified by this resource. See the following sub-section for an example.

5.16.4 Required JSON Format Resource

The required JSON *format resource* is defined as:

```
{
  "uri": URI,
  "name": "JSON",
  "description": "JavaScript Object Notation",
  "tags": String[] ?,
  "mime_type": "application/json",
  "version": "RFC4627",
  "documentation": "http://www.ietf.org/rfc/rfc4627.txt",
  "metadata": object
}
```

The *name*, *mime_type*, *version*, and *documentation* attribute values for the JSON *format resource* SHALL reflect the above values. [RE-42]

5.17 type_definition Resource

A *type_definition resource* describes a resource type supported by the Platform. This resource is a sub-type of the *collection resource* and has the following pseudo schema:

```
{
  "uri": URI,
  "name": String,
  "description": String ?,
  "tags": String[] ?,
  "representation_skew": String ?,
  "external_management_resource": URI ?,
  "collection_type": URI,
  "total_items": Number,
  "items_per_page": Number,
  "start_index": Number,
  "items": attribute_definition[],
  "documentation": URI,
  "inherits_from_collection": URI ?,
  "metadata": object
}
```

The *type_definition resource* contains the following attributes:

5.17.1 items

Type: `attribute_definition[]`

Required: true

The value of this attribute is an array of *attribute_definition resources* that enumerate the attributes of the type described by this resource.

For every attribute of the type that is not inherited from its super-types, the `items` array SHALL contain an *attribute_definition resource* that defines that attribute. [RE-45] In cases where a sub-type adds additional constraints to an attribute inherited from its super-types (e.g. makes an optional attribute required), Providers SHALL include an *attribute_definition resource* for that attribute in the `items` array. [RE-76] For more information on the *attribute_definition resource* see the next section.

5.17.2 documentation

Type: URI

Required: true

This attribute contains a URI that points to the documentation for the resource type. For resource types that are defined in this specification, the URI can point to this specification.

5.17.3 inherits_from_collection

Type: URI

Required: false

This attribute references a collection of *type_definition resources* that enumerate the type(s) that the described type inherits from. Items in this collection SHALL NOT, either directly or transitively, refer to the described resource. [MO-06] If a type inherits only from the *camp_resource* type then this attribute MAY be absent. [MO-07]

5.18 attribute_definition Resource

An *attribute_definition resource* represents exactly one supported attribute of one or more resource types. This resource has the following pseudo schema:

```
{
  "uri": URI,
  "name": String,
  "description": String ?,
  "tags": String[] ?,
  "representation_skew": String ?,
  "external_management_resource": URI ?,
  "documentation": URI,
  "attribute_type": String
  "required": Boolean,
  "metadata": object
}
```

Instances of the *attribute_definition resource* contain the following attributes:

5.18.1 documentation

Type: URI

Required: true

The value of the `documentation` attribute is a URI that references the documentation for the attribute that this resource represents. For attributes that are defined in this specification, this URI references this specification.

5.18.2 attribute_type

Type: String

Required: true

Mutable: false

The value of the `attribute_type` attribute specifies the type of the attribute that is described by this resource. See Section 5.2, “[Attribute Types](#)”, for a list of the values defined by this specification.

The appearance of the square bracket symbols, “[]”, appended to the value of the `attribute_type` attribute indicates that the value of the attribute that is described by this resource is an array of the specified type. For example, an `attribute_type` value of “String[]” indicates that the value of the attribute being described by is an array of Strings.

5.18.3 required

Type: Boolean

Required: true

Mutable: false

The value of the `required` attribute determines if the attribute defined this *attribute_definition resource* is required for resources of the type defined by the containing *type_definition resource*. A value of “true” indicates that the referenced attribute will always be present in resources of the type defined by the containing *type_definition resource*.

5.19 parameter_definition Resource

A *parameter_definition resource* represents exactly one supported parameter of one or more resource types. This resource has the following pseudo schema:

```
{
  "uri": URI,
  "name": String,
  "description": String ?,
  "tags": String[] ?,
  "representation_skew": String ?,
  "external_management_resource": URI ?,
  "parameter_type": String,
  "required": Boolean,
  "default_value": <parameter type> ?,
  "parameter_extension": URI ?,
  "metadata": object
}
```

parameter_definition resources contain the following attributes:

5.19.1 parameter_type

Type: String

Required: true

This attribute specifies the type of the attribute that this resource represents. For example, “String”, “Timestamp”.

5.19.2 required

Type: Boolean

Required: true

The value of the `required` attribute specifies whether the parameter defined by this *parameter_definition resource* is required for HTTP POST requests on the resource that references the collection that contains

this *parameter_definition* resource. A value of “true” indicates that the referenced parameter is required for all POST requests.

5.19.3 default_value

Type: As defined by *parameter_definition* resource.

Required: false

The value of the `default_value` attribute, when present, specifies the default value for the parameter defined by this *parameter_definition* resource. If the Consumer does not supply a value for the parameter, the value of this attribute will be used. Note that the presence of the `default_value` attribute is mutually exclusive with a required value (see above) of “true”.

5.19.4 parameter_extension

Type: URI

Required: false

If this parameter is handled by an extension, this attribute refers to the *extension* resource that represents that Extension and documents how the parameter is handled.

5.20 operation Resource

An *operation* resource represents exactly one operation or action available on a target resource. This resource has the following pseudo schema:

```
{
  "uri": URI,
  "name": String,
  "description": String ?,
  "tags": String[] ?,
  "representation_skew": String ?,
  "external_management_resource": URI ?,
  "documentation": URI ?,
  "target_resource": URI,
  "parameter_definition_collection": URI ?,
  "metadata": object
}
```

In addition to the methods defined in Section 5.5, “HTTP Method Support”, Providers SHALL support the HTTP POST method on the *operation* resource. [RE-64]

A POST request on the *operation* resource invokes the Operation on the target resource. The Operation MAY require content in the body of the POST, such as parameters. [RE-47] The response to a POST request on an *operation* resource SHOULD indicate what changes were made on the target resource. [RE-48] For asynchronous operations, the response SHOULD indicate how to track the progress of the request operation. [RE-49]

NOTE: CAMP does not specify a mechanism for asynchronous operations. Consumers and Providers can use mechanisms available to them for asynchronous operations, for example, HTTP callbacks. Consumers can also poll the target resource and use the value of the `representation_skew` attribute (see Section 5.4.5, “`representation_skew`”) and/or other attributes as an indicator of any changes in that resource.

Instances of the *operation* resource contain the following attributes:

5.20.1 name

Type: String

Required: true

This attribute contains the name of the operation that this resource represents. For example, “deploy” or “resize”.

5.20.2 documentation

Type: URI

Required: false

If present, the value of this attribute is a URI that refers to the documentation for the operation represented by this resource.

5.20.3 target_resource

Type: URI

Required: true

This attribute indicates the CAMP resource on which the linked operation is invoked.

5.20.4 parameter_definition_collection

Type: URI

Required: false

The value of the `parameter_definition_collection` attribute references a collection of *parameter_definition resources* that describe the parameters accepted by this operation on an HTTP POST request. Each of the *parameter_definition resources* provides metadata for a parameter as described in Section 5.19, “[parameter_definition Resource](#)”.

This attribute can be omitted when the operation does not support parameterization.

5.21 sensor Resource

A *sensor resource* represents exactly one supported sensor on one or more resources. A *sensor resource* represents dynamic data about resources, such as metrics or state. A *sensor resource* is useful for exposing data that changes rapidly, or that may need to be fetched from a secondary system. This resource has the following pseudo schema:

```
{
  "uri": URI,
  "name": String,
  "description": String ?,
  "tags": String[] ?,
  "representation_skew": String ?,
  "external_management_resource": URI ?,
  "documentation": URI ?,
  "target_resource": URI,
  "sensor_type": String | URI,
  "value": <sensor_type> ?,
  "units": String ?,
  "timestamp": Timestamp ?,
  "operation_collection": URI ?,
  "metadata": object
}
```

Instances of the *sensor resource* contain the following attributes:

5.21.1 documentation

Type: URI

Required: false

If present, the value of this attribute is a URI that refers to the documentation for the sensor represented by this resource.

5.21.2 target_resource

Type: URI

Required: true

This attribute indicates the CAMP resource for which this *sensor resource* supplies runtime data.

5.21.3 sensor_type

Type: String | URI

Required: true

This attribute specifies the type of the data that this *sensor resource* collects. If the value of the “value” attribute is a CAMP-defined scalar, the value of this attribute will be as defined in Section 5.2, “[Attribute Types](#)”; for example, “String”, “Number”, etc. If the value of the “value” attribute is an object, the value of this attribute will be a URI that references the *type_definition resource* that defines that object. See Section 5.17, “[type_definition Resource](#)”.

5.21.4 value

Type: As defined by sensor_type

Required: false

This attribute contains the current or most recent available value for this sensor. It can be omitted, for example, to indicate that no current value is available; either because no data has been collected or the collected data is stale.

5.21.5 units

Type: String

Required: false

The value of this optional attribute identifies the unit of measure associated with the value of the value attribute.

5.21.6 timestamp

Type: Timestamp

Required: false

This attribute contains the timestamp of the last collection or relevant activity of the sensor. When a “value” attribute is supplied, any timestamp provided in this attribute SHOULD correspond to when that value was observed. [\[RE-51\]](#)

5.21.7 operation_collection

Type: URI

Required: false

This attribute references a collection of *operation resources*. The *operation resources* in this collection represent the operations available for this sensor.

Extensions MAY be defined to govern common sensor management operations, such as enabling, disabling, adjusting collection frequency, specifying the history of values which should be remembered, or collecting immediately. [\[RE-52\]](#)

6 Protocol

6.1 Transfer Protocol

The CAMP API is based on the Hypertext Transfer Protocol, version 1.1 [RFC2616]. Requests sent from Consumers across unsecured networks SHOULD use the HTTPS protocol. [PR-40] TLS 1.1 [RFC4346] SHALL be implemented by the Provider. [PR-41] TLS 1.2 [RFC5246] is RECOMMENDED. [PR-42] When TLS is implemented, the following cipher suites are RECOMMENDED to ensure a minimum level of security and interoperability between implementations:

- TLS_RSA_WITH_AES_128_CBC_SHA (mandatory for TLS 1.1/1.2) [PR-43]
- TLS_RSA_WITH_AES_256_CBC_SHA256 (addresses 112-bit security strength requirements) [PR-44]

Note: For interoperability reasons, Providers are encouraged to support a common authentication scheme in order to simplify the implementation of client tools that are intended to work with multiple Providers. The *platform_endpoint_resource* `auth_scheme` attribute (see Section 5.8.6, “`auth_scheme`”) makes available authentication schemes discoverable by unauthenticated clients.

6.2 URI Space

The resources in the system are identified by URIs. Dereferencing the URI will yield a representation of the resource containing resource attributes and links to associated resources.

Note: Consumers are cautioned against making assumptions about the layout of the URIs or the structure of the URIs of the resources.

6.3 Media Types

6.3.1 Required Formats

In this specification resource representations are encoded in JSON, as specified in [RFC4627]. The media-type associated with CAMP JSON resource representations is “application/json”.

Providers SHALL provide representations of all available resources in JSON. [PR-01]

6.3.1.1 Duplicate Keys in JSON Objects

CAMP defined JSON objects do not contain duplicate keys. Consumers and Providers SHALL NOT transmit JSON objects that contain duplicate keys. [PR-02]

If a Consumer sends a Provider a request containing duplicate keys in a JSON object, the Provider SHOULD reject the request by sending back a ‘400 Bad Request’ status code. [PR-03] If a Provider sends a Consumer a response containing duplicate keys in a JSON object, the Consumer SHOULD raise an error to the user indicating the response from the server was malformed. [PR-04]

Note: Duplicate keys in JSON objects are allowed by JSON [RFC4627]. This specification prohibits duplicate keys for interoperability reasons.

6.3.2 Supported Formats

If Supported Formats besides JSON are defined in the *format_resource* collection referenced by the `supported_format_collection` attribute of the *platform_resource*, then the indicated resource representations are allowed in the Supported Formats.

For each Supported Format, Consumers MAY request any resource from the Provider in that format. [PR-45] Providers SHALL respond in the requested Supported Format. [PR-05]

A client can request any Supported Format using HTTP content negotiation.

6.4 Request Headers

All PUT requests that update a resource SHOULD contain the If-Match header field with a single entity tag value. [PR-06] If the If-Match header field value in the request does not match the one on the server-side, the Provider SHALL send back a '412 Precondition Failed' status code. [PR-07]

6.5 Query String Parameters

The CAMP API supports a set of request parameters that are conveyed in the query component of the URI. The following table lists the supported key values for these parameters.

Key	Description	Example
select_attr	Comma (“,”) separated attribute names of the resource to return. See Section 7.3.1, “ Selecting Attributes on a GET ”.	assembly132?select_attr=name%2Cdescription%2Ctags Would access only “name”, “description”, “tags” attributes of assembly132.
select_collection_attr	Comma (“,”) separated list for the attribute names of the resources in the collection to return. See Section 7.3.2, “ Selecting Attributes of the Resources in a Collection on a GET ”.	If /paas/assembly/8675309/comps is the components resource URI of assembly/8675309, then GET /paas/assembly/8675309/comps?select_collection_attr=name%2Cdescription%2Ctags would access only “name”, “description”, “tags” attributes of the components of assembly/8675309.
sort	The value of this parameter is a sequence of attribute names separated by commas with each attribute prefixed by the symbol ‘+’ implying ascending, or ‘-’ implying descending sort order. See Section 7.3.3, “ Sorting Elements of a Collection ” for details.	assembly_factory?sort=+name Would sort the items array of the assembly factory resource by the name attribute of the elements in that array in ascending order.
start_index	Used to specify the first requested element in a <i>collection resource</i> . See Section 7.3.4 “ Pagination of a Collection ” for details.	assembly_factory?start_index=10 Would retrieve the elements of the <i>assembly factory</i> resource starting from the tenth element in the collection.
max_page	Used to page through the elements in a <i>collection resource</i> . See Section 7.3.4 “ Pagination of a Collection ” for details.	assembly_factory?max_page=10 Would retrieve the first 10 elements of the <i>assembly factory</i> resource.
index_in_collection	Used to get the index of a specified resource within a collection. See Section 7.3.5 “ Retrieving Index of an Item in a Collection ” for details.	assemblies?index_in_collection=http%3A%2F%2Fexample.org%2Fmy_paas%2Fapps%2F1 Would retrieve a collection with a single element (the one identified by “http://example.org/my_paas/apps/1”). The attributes of the <i>collection resource</i> indicate where in the collection the supplied element exists.

Table 6-1: Query String Parameters

6.6 POST Body Parameters

Parameters MAY be included when performing a POST request on any resource with a parameter_definition_collection attribute defined. [PR-14] Supported parameters are defined by the collection of *parameter_definition resources* referenced by the parameter_definition_collection attribute of the resource handling the POST request.

Example of a POST Parameter:

```
POST /<assembly-template-resource-url> HTTP/1.1
Host: example.org
Content-Type: application/json
Content-Length: ...

{ "EXAMPLE:someParameter": "bar" }

HTTP/1.1 201 Created
Location: http://example.org/paas/assembly/1
Content-Type: ...
Content-Length: ...
```

6.6.1 Parameter Handling

Parameters allow customizing resources upon creation. Parameters MAY have the same name as an attribute of the resource. [PR-15] In such cases the Provider SHOULD set the value of that attribute to the value of the parameter OR clearly document alternate behavior. [PR-16] The parameter_extension attribute MAY be used to reference the extension which documents how the parameter is handled. [PR-17]

If a POST request body does not contain a value for a required parameter, a “400 Bad Request” response SHALL be returned. [PR-18]

If a POST request body does not contain an acceptable value for a parameter, a “400 Bad Request” response SHALL be returned. [PR-19]

6.7 Response Headers

Responses returned by the Provider make standard use of HTTP headers. All HTTP responses that return representation of a resource SHOULD use a strong ETag response header field indicating the current value of the entity tag for the resource. [PR-20]

6.8 HTTP Status Codes

The API returns standard HTTP response codes.

7 Operations

7.1 Deploying an Application

Deploying an application uploads the artifacts and metadata that make up the application, allocates the necessary resources and services, and, in the successful case, creates a running application (represented by an *assembly resource*).

There are two ways to deploy an application using a PDP: by POSTing the entire PDP to the *assembly_factory resource* (by value) or by POSTing the URI of the PDP to the *assembly_factory resource* (by reference). Similarly, there are two ways to deploy an application using a Plan: by POSTing the entire Plan file to the *assembly_factory resource* (by value) or by POSTing the URI of the Plan file or *plan resource* to the *assembly_factory resource* (by reference). All of these methods are described below. Providers SHALL support PDPs that use either the ZIP [ZIP], TAR [TAR], or GZIP [RFC1952] compressed TAR formats. [RMR-13]

7.1.1 Deploying an Application by Reference

To deploy an application by reference, a Consumer sends a reference to either a PDP, Plan file, or *plan resource* in a POST request to the *assembly_factory resource* (see Section 5.10, “[assembly_factory Resource](#)”). Providers SHALL support the deployment of applications via HTTP POST requests on the *assembly_factory resource* as described in this section. [PR-49] The entity body of the request contains a URI that identifies the PDP, Plan file, or *plan resource* that is being deployed. If the URI that identifies the PDP, Plan file, or *plan resource* is a relative URI, its base URI is that of the *platform resource*. When deploying a PDP the JSON serialization of the HTTP request entity body is:

```
{"pdp_uri": "<uri-of-the-pdp>"}
```

When deploying a Plan file or *plan_resource* the JSON serialization of the HTTP request entity body is:

```
{"plan_uri": "<uri-of-the-plan>"}
```

Where, the value of *pdp_uri* is the URI of the PDP to be deployed and the value of *plan_uri* is the URI of the Plan to be deployed. To support the deployment of applications via a reference to either a PDP, Plan file, or *plan resource*, Providers SHALL accept the "application/json" media type. [PR-68] The JSON object MAY contain additional name-value pairs that are not defined in this specification. [PR-33]

Note that the value of *plan_uri* can refer to either a Plan file or a *plan resource*. A referenced *plan resource* can exist either on the same CAMP Platform as the target *assembly_factory resource*, or on some other CAMP Platform. In the case where the referenced *plan resource* exists within the same Platform as the target *assembly_factory resource*, Consumers are advised to use a relative URL for the *plan resource* reference to help Providers identify the *plan resource* as local.

An example HTTP request-response is as follows:

```
POST /paas/assemblies HTTP/1.1
Host: example.org
Content-Type: application/json
Content-Length: 46
...
{"pdp_uri": "/paas/pdp/1",
 "description": "Mike's other Drupal instance"}
HTTP/1.1 201 Created
Location: http://example.org/paas/assembly/11
...
```

Note the inclusion of *description* parameter/attribute in the body of the above POST request.

On successfully processing the request the Provider SHALL create an *assembly resource* and return a 201 Created status code in the HTTP response. [PR-50] The Provider SHALL include the *Location* header in the HTTP response and the value of this header SHALL reference the newly created *assembly resource*. [PR-51] The Provider SHALL update the *items* attribute of the *assembly_factory resource* to include the newly created *assembly resource*. [PR-52]

7.1.2 Deploying an Application by Value

To deploy an application by value, a Consumer transmits the contents of either a PDP or a Plan file in a POST request to the top-level *assembly_factory resource* (see Section 5.10, “*assembly_factory Resource*”). The Consumer can do this in two ways, either by including the data as a part in a multipart MIME message or by simply including the data in the entity body of the HTTP request.

To support the deployment of applications using a PDP, Providers SHALL accept the media types associated with the various formats as follows:

- ZIP: "application/x-zip" [PR-29]
- TAR: "application/x-tar" [PR-30]
- GZIP compressed TAR: "application/x-tgz" [PR-31]

To support the deployment of applications using a Plan file, Providers SHALL accept the use of the "application/x-yaml" media type. [PR-32]

On successfully processing the request the Provider SHALL create an *assembly resource* and return a 201 *Created* status code in the HTTP response. [PR-53] The Provider SHALL include the *Location* header in the HTTP response and the value of this header SHALL reference the newly created *assembly resource*. [PR-54] The Provider SHALL update the *items* attribute of the *assembly_factory resource* to include the newly created *assembly resource*. [PR-55]

For large PDPs, the Consumer can use existing HTTP facilities like chunked transfer encoding.

7.1.2.1 Deploying an Application by Value Using MIME

Providers SHALL support the deployment of applications via HTTP POST requests on the *assembly_factory resource* as described in this section. [PR-74] The entity body of the request is a multipart MIME message that contains the PDP or the Plan file that is being deployed. The value of the HTTP Content-Type header is "multipart/form-data" [RFC2388]. The MIME part that contains the file data has the value of the name parameter of its Content-Disposition header set to "pdp_file", if a PDP is being deployed, or "plan_file", if a Plan file is being deployed. CAMP parameters can be included in the request as additional MIME parts using the value of the name parameter of the Content-Disposition header to indicate the CAMP parameter being included.

An example HTTP request-response is as follows:

```
POST /paas/assemblies HTTP/1.1
Host: example.org
Content-Length: 9768506
Content-Type: multipart/form-data; boundary=----WebKitFormBoundaryU6AnBoMx
...
-----WebKitFormBoundaryU6AnBoMx
Content-Disposition: form-data; name="pdp_file"; filename="Mike_Drupal2.pdp"
Content-Type: application/x-zip

... binary octets ...
-----WebKitFormBoundaryU6AnBoMx
Content-Disposition: form-data; name="description"

Mike's other Drupal instance
-----WebKitFormBoundaryU6AnBoMx--
```

```
HTTP/1.1 201 Created
Location: http://example.org/paas/assembly/12
Content-Type: application/json
...
```

Note the inclusion of the description parameter as a separate MIME part.

7.1.2.2 Deploying an Application by Value Without MIME

Providers SHALL support the deployment of applications via HTTP POST requests on the *assembly_factory resource* in which the entity body of the request contains the PDP or the Plan file that is being deployed. [PR-60]

An example HTTP request-response is as follows:

```
POST /paas/assemblies HTTP/1.1
Host: example.org
Content-Length: 976361
Content-Type: application/x-zip
Content-Transfer-Encoding: binary
...
... binary PDP ZIP octets ...

HTTP/1.1 201 Created
Location: http://example.org/paas/assembly/12
...
```

Note that it is not possible to include additional parameters when using this mechanism to deploy by an application.

7.2 Registering a Plan

As described in Section 2.3, “Deployment”, CAMP implementations can choose to support the expression of Plans as CAMP resources. This feature allows Consumers to register an application (upload the artifacts and metadata, validate the Plan, resolve dependencies, etc.) without creating a running instance of that application. Consumers can later instantiate an application from the *plan resource*.

Similarly to deploying an application, Plans can be registered using either a PDP or a Plan file. The PDP or Plan file can be supplied by value or by reference.

The archive formats available to the PDP are identical to those used to deploy an application.

7.2.1 Registering a Plan by Reference

To register a Plan by reference, a Consumer sends a reference to either a PDP or a Plan file in a POST request to the *plan resource* collection. Providers that support the *plan resource* and *plan resource* collection SHALL support the registration of Plans via an HTTP POST request on the *plan resource* collection as described in this section. [PR-56] The entity body of the request contains a URI that identifies the PDP or the Plan file that is being registered. If the URI that identifies the PDP or the Plan file is a relative URI, its base URI is that of the *platform resource*. When registering a PDP the JSON serialization of the HTTP request entity body is:

```
{"pdp_uri": "<uri-of-the-pdp>"}
```

When registering a Plan file the JSON serialization of the HTTP request entity body is:

```
{"plan_uri": "<uri-of-the-plan>"}
```

Where, the value of `pdp_uri` is the URI of the PDP to be registered and the value of `plan_uri` is the URI of the Plan to be registered. To support the registration of Plans via a reference to either a PDP or a Plan file, Providers SHALL accept the "application/json" media type. [PR-69] The JSON object MAY contain additional name-value pairs that are not defined in this specification. [PR-46]

An example HTTP request-response is as follows:

```
POST /paas/plans HTTP/1.1
Host: example.org
Content-Type: application/json
Content-Length: 46
...

{"pdp_uri": "/paas/pdp/1",
 "description": "Mike's other Drupal instance"}

HTTP/1.1 201 Created
Location: http://example.org/paas/plan/9
...
```

Note the inclusion of description parameter/attribute in the body of the above POST request.

On successfully processing the request the Provider SHALL create a *plan resource* and return a 201 Created status code in the HTTP response. [PR-57] The Provider SHALL include the *Location* header in the HTTP response and the value of this header SHALL reference the newly created *plan resource*. [PR-58] The Provider SHALL update the *items* attribute of the *plan resource* collection to include the newly created *plan resource*. [PR-59]

7.2.2 Registering a Plan by Value

To register a Plan by value, a Consumer transmits the contents of either a PDP or a Plan file in a POST request to the *plan resource* collection. The Consumer can do this in two ways, either by including the data as a part in a multipart MIME message or by simply including the data in the entity body of the HTTP request.

To support the registration of Plans using a PDP, Providers SHALL accept the media types associated with the various formats as follows:

- ZIP: "application/x-zip" [PR-70]
- TAR: "application/x-tar" [PR-71]
- GZIP compressed TAR: "application/x-tgz" [PR-72]

To support the registration of Plans using a Plan file, Providers SHALL accept the use of the "application/x-yaml" media type. [PR-73]

On successfully processing the request the Provider SHALL create a *plan resource* and return a 201 Created status code in the HTTP response. [PR-62] The Provider SHALL include the *Location* header in the HTTP response and the value of this header SHALL reference the newly created *plan resource*. [PR-63] The Provider SHALL update the *items* attribute of the *plan resource* collection to include the newly created *plan resource*. [PR-64]

For large PDPs, the Consumer can use existing HTTP facilities like chunked transfer encoding.

7.2.2.1 Registering a Plan by Value Using MIME

Providers that support *plan resources* and the *plan resource* collection SHALL support the registration of Plans via HTTP POST requests on the *plan resource* collection as described in this section. [PR-75] The entity body of the request is a multipart MIME message that contains the PDP or the Plan file that is being registered. The value of the HTTP Content-Type header is "multipart/form-data" [RFC2388]. The MIME part that contains the file data has the value of the name parameter of its Content-Disposition header set to "pdp_file", if a PDP is being registered, or "plan_file", if a Plan file is being registered. CAMP parameters can be included in the request as additional MIME parts using the value of the name parameter of the Content-Disposition header to indicate the CAMP parameter being included.

An example HTTP request-response is as follows:

```

POST /paas/plans HTTP/1.1
Host: example.org
Content-Length: 1685
Content-Type: multipart/form-data; boundary=----WebKitFormBoundaryE733b300
...

-----WebKitFormBoundaryE733b300
Content-Disposition: form-data; name="plan_file"; filename="Mike_Drupal.yaml"
Content-Type: application/x-yaml

... unicode characters ...

-----WebKitFormBoundaryE733b300
Content-Disposition: form-data; name="description"

Mike's other Drupal instance

-----WebKitFormBoundaryE733b300--

HTTP/1.1 201 Created
Location: http://example.org/paas/plan/9
...

```

Note the inclusion of the description parameter as a separate MIME part.

7.2.2.2 Registering a Plan by Value Without MIME

Providers that support *plan resources* and the *plan resource* collection SHALL support the registration of Plans via HTTP POST requests on the *plan resource* collection in which the entity body of the request contains the PDP or the Plan file that is being registered. [\[PR-61\]](#)

An example HTTP request-response is as follows:

```

POST /paas/plans HTTP/1.1
Host: example.org
Content-Length: 1465
Content-Type: application/x-yaml
...

... unicode characters ...

HTTP/1.1 201 Created
Location: http://example.org/paas/plan/9
...

```

Note that it is not possible to include additional parameters when using this mechanism to register a Plan.

7.3 Retrieving Resources

7.3.1 Selecting Attributes on a GET

To retrieve a subset of the attributes in a resource, the Consumer MAY use the 'select_attr' query parameter in conjunction with the HTTP GET method. [\[PR-08\]](#) A Provider SHALL return only those attributes of the queried resource whose name occurs in the list specified by the value of 'select_attr'. [\[PR-47\]](#) If an attribute listed in the value of the 'select_attr' query parameter is not part of the resource, a "400 Bad Request" status code SHALL be returned. [\[PR-09\]](#) The "select_attr" query parameter MAY appear more than once (separated by an "&"). [\[PR-10\]](#)

7.3.2 Selecting Attributes of the Resources in a Collection on a GET

To retrieve a subset of the attributes of the resources in a collection, the Consumer MAY use the 'select_collection_attr' query parameter in conjunction with the HTTP GET method. [\[PR-77\]](#) A Provider SHALL return only those attributes of the resources in the collection whose name occurs in the list specified by the value of 'select_collection_attr'. [\[PR-78\]](#) The absence of one or more attributes listed in

'select_collection_attr' in one or more items in the collection is not considered an error. If an attribute listed in the value of the 'select_collection_attr' query parameter is not part of the resource, it SHALL not appear in the returned resource. [PR-79] If a resource in the collection contains none of the attributes listed in the 'select_collection_attr' query parameter, then it SHALL be returned as an empty object. [PR-80] The 'select_collection_attr' query parameter MAY appear more than once (separated by an "&"), in which case the value of the query parameter is considered to be the union of the attribute sets specified by each individual query parameter. [PR-81]

If the 'select_collection_attr' query parameter is supplied for a URL that is not a collection resource, the provider SHALL return 400 Bad Request. [PR-82]

7.3.2.1 Uniqueness of Resources in a Collection

In the case where no 'select_collection_attr' parameter is specified, all the elements of a collection are unique by virtue of the value of their uri attribute. However, when the 'select_collection_attr' query parameter is applied to the attributes of these resources, elements of the resulting collection could contain duplicate attribute sets. Providers SHALL remove all duplicates from *collection resources* and return a unique collection of sets of attributes. [PR-83]. In cases where pagination is used in combination with 'select_collection_attr', the Provider SHALL apply pagination to the unique collection of sets of attributes. [PR-84]. Note that a corollary to the previous requirement is that the values for the total_items, items_per_page, and start_index attributes all are calculated with respect to the collection after any duplicate elements have been removed (i.e. the "unique collection").

For example, consider the following collection of *component resources*:

```
{
  "uri": "http://example.org/botsc/assembly/8675309/comps",
  "name": "BOTS-Test-App-3 Components",
  ...
  "total_items": 3,
  "items_per_page": 3,
  "start_index": 0,
  "items" : [
    {
      "uri": "http://example.org/botsc/comp/7429",
      "name": "MySQL5-med-20160821",
      "description": "MySQL5.7 on medium VM",
      "status": "RUNNING",
      ...
    },
    {
      "uri": "http://example.org/botsc/comp/275831/comps",
      "name": "Kafka9-20161003"
      "description": "Shared Kafka 0.9.0.1",
      "status": "PAUSED",
      ...
    },
    {
      "uri": "http://example.org/botsc/comp/275831/comps",
      "name": "Logstash-20170112"
      "status": "RUNNING",
      ...
    }
  ]
}
```

The following request

```
GET /botsc/assembly/8675309/comps?select_collection_attr=description HTTP/1.1
Host: example.org
...
```

will return the following resource. Note how the lack of a description attribute for the Logstash component resulted in an empty element in the items array.

```

{
  "uri": "http://example.org/botsc/assembly/8675309/comps",
  "name": "BOTS-Test-App-3 Components",
  ...
  "total_items": 3,
  "items_per_page": 3,
  "start_index": 0,
  "items" : [
    {
      "description": "MySQL5.7 on medium VM",
    },
    {
      "description": "Shared Kafka 0.9.0.1",
    },
    {
    }
  ]
}

```

The following request on the same resource

```

GET /botsc/assembly/8675309/comps?select_collection_attr=status HTTP/1.1
Host: example.org
...

```

will produce the following resource. Note how the `total_items` attribute represents the total number resources after duplicate elimination.

```

{
  "uri": "http://example.org/botsc/assembly/8675309/comps",
  "name": "BOTS-Test-App-3 Components",
  ...
  "total_items": 2,
  "items_per_page": 2,
  "start_index": 0,
  "items" : [
    {
      "status": "RUNNING",
    },
    {
      "status": "PAUSED",
    }
  ]
}

```

7.3.3 Sorting Elements of a Collection

The Consumer MAY specify the 'sort' query parameter when sending an HTTP GET request to a *collection resource*. [OP-01] The value of the 'sort' query parameter is a sequence of comma-separated attribute names, with each name prefixed by the '+' or '-' characters. The prefix indicates the sort order for the attribute that it prefixes; the '+' indicates ascending sort order and the '-' indicates descending sort order. If neither the '+' nor the '-' prefix is present, then '+' is assumed. The order of the attribute names within the sequence indicates their precedence in the sort operation. Names that appear earlier in the sequence have precedence over the names that appear later in the sequence.

For example, the results of a GET request on the component collection of an *assembly resource* can be sorted with the following query parameter:

```

GET /paas/assembly/8675309/comps?sort=-status,+name

```

This specifies that the `items` array of the returned *collection resource* should first be sorted in descending order by the value of the `status` attribute and then, within components which share the same status value, in ascending order by the value of the `name` attribute.

The Provider SHALL use the value of the 'sort' query parameter to order the elements of the `items` array in the returned *collection resource*. [OP-02] In cases where pagination is used in combination with sorting, the Provider SHALL perform the sort operation before applying pagination. [OP-03]

7.3.3.1 Sort Ordering

The sort operation is defined for attributes of the resources within the collection (i.e. the elements of the `items` array), not for the attributes of the *collection resource* itself. It applies only to the scalar CAMP types (see Section 5.2, "Attribute Types"). Note that it is possible that the sorted attribute(s) may not be present in all elements of the collection, in which case the non-existent attributes are treated as having the value `null`. For all scalar attribute types a value of `null` is considered to be the lowest value; appearing first when sorting in ascending order and last when sorting in descending order. A Provider SHALL respond to a request in which the value of the 'sort' query parameter specifies one or more non-scalar attributes (e.g. an attribute whose value is an array) with a "400 Bad Request" status code. [OP-04]

7.3.3.1.1 Boolean

`false` has a lower value than `true`.

7.3.3.1.2 String

Based on Unicode collation algorithm [UTS10].

7.3.3.1.3 Number

In the numeric order of the encoded number (as per JSON [RFC4627]).

7.3.3.1.4 URI

As per String.

7.3.3.1.5 Timestamp

In the chronological order of the date corresponding to the string (expressed as per ISO 8601 [ISO 8601:2004]).

7.3.4 Pagination of a Collection

Some collections may contain too many elements to be conveniently and/or efficiently handled by either the Consumer or the Provider. As described in Section 5.6, "collection Resource", the elements returned in a representation of a *collection resource* are a subset (possibly equal to the entire set) of the collection. Consumers can deliberately select specific subsets of the collection through the use of the `start_index` and `max_page` query parameters

A Consumer MAY specify the 'start_index' and 'max_page' query parameters when sending an HTTP GET request to a *collection resource*. [OP-05] The `start_index` query parameter specifies the zero-based index of the first element in the subset of the collection that the Consumer wishes to retrieve. The `max_page` query parameter specifies the maximum number of elements the Consumer wishes to retrieve. The absence of the `start_index` query parameter is treated as a request with a `start_index` of 0. The absence of a `max_page` query parameter is treated as a request with a `max_page` of ∞ . If provided, the value of the 'start_index' query parameter SHALL be a non-negative integer. [OP-06] If provided, the value of the 'max_page' query parameter SHALL be a positive integer. [OP-07]

For example, the following sequence of requests would page through elements in a collection resource at ten items per page:

```
GET /my_paas/assembly_factory?max_page=10 HTTP/1.1
GET /my_paas/assembly_factory?start_index=10&max_page=10 HTTP/1.1
```

```
GET /my_paas/assembly_factory?start_index=20&max_page=10 HTTP/1.1
```

Because pagination is not stateful, Consumers need to be prepared to handle inconsistent results. For example suppose that, between the second and third requests in the above example, a new *assembly resource* was created and inserted at index 17 in the *assembly factory*. This would cause the first element of the *collection resource* retrieved from the third request to be a repeat of the last element of the *collection resource* retrieved from the second request. There a number of strategies for detecting and dealing with these situations. If the Provider returns ETag headers, the Consumer can use the value of this header to detect when items have been added to or removed from the collection. Furthermore, because the elements of the *items* array are CAMP resources and all CAMP resources are uniquely identified by the value of their *uri* attribute, overlaps like the one discussed above can be detected by the appearance of one or more resources in two separate “pages” in a series of requests. The possibility of missing items (due to deletion of previously seen items in the collection, the opposite of the above example) can be mitigated (but not completely removed) by purposely overlapping pages and correcting for the overlap by comparing the values of the *uri* attributes of the elements in the *items* array.

A Provider MAY respond to an HTTP GET request on a collection resource with a representation whose *items_per_page* value is less than the number of items requested by the Consumer (whether explicitly via the ‘*max_page*’ query parameter or implicitly through the absence of this query parameter). [OP-08] A Provider SHALL NOT respond to an HTTP GET request on a collection resource with a representation whose *items_per_page* value is greater than the value of the ‘*max_page*’ query parameter. [OP-09] A Provider SHALL respond to a request in which the value of the ‘*start_index*’ query parameter is equal to or greater than the number of items in the collection with a “400 Bad Request” status code. [OP-10]

7.3.5 Retrieving Index of an Item in a Collection

In many user interface use cases it is useful to find the index of an item within a collection. To retrieve the index of an item within a collection, the Consumer MAY use the *index_in_collection* query parameter in conjunction with the HTTP GET method. [OP-11] The value of this query parameter is the URI of the resource of interest within the collection. If the specified resource is not present in the collection a “404 Not Found” status code SHALL be returned by the Provider. [OP-12]

If the specified resource is present in the collection the Provider will return a *collection resource* with the following constraints:

- The *items_per_page* attribute SHALL be 1. [OP-13]
- The *items* array SHALL contain the resource referenced by value of the query parameter. [OP-14]

The normal semantics of a *collection resource* provides the Consumer with enough information to determine where the item falls within the collection. The value of the *total_items* attribute is the total number of resources in the collection and the value of the *start_index* attribute is the index of the resource of interest.

7.4 Updating Resources

Attributes that are referenced by a value in the *consumer_mutable* array can be modified by Consumers in two ways. Consumers MAY use the HTTP PUT method to replace the representation of a resource, in its entirety, with a new representation that adds, omits or replaces the values for some of the attributes. [PR-23] Alternatively, Consumers MAY use the HTTP PATCH [RFC5789] method and the “application/json-patch+json” media type [RFC6902] to add, delete, or replace specific attributes. [PR-24]

See Section 5.5, “HTTP Method Support”, for a discussion of the interaction of the POST, PUT, and PATCH methods with mutable and consumer-mutable attributes.

7.4.1 Updating with PUT

HTTP PUT requests are requests for complete replacement of the resource identified by the request URL.

On successfully processing an HTTP PUT request a Provider SHALL update all the consumer-mutable attributes of the target resource, and only these, with the values of the matching attributes in the request.

[PR-48] If a resource attribute is present on a resource and if an HTTP PUT request omits that attribute, it SHOULD be treated by the Provider as a request to delete the attribute. [PR-25]

7.4.1.1 Partial Updates with PUT

Section 6.5, “[Query String Parameters](#)”, describes the use of the ‘select_attr’ query parameter to circumscribe the GET method to a subset of a resource’s attributes. To allow an update of a subset of a resource’s attributes, Providers SHALL support the use of the ‘select_attr’ query parameter in conjunction with the HTTP PUT method. [PR-76] A Consumer SHALL NOT include attributes, whose name does not occur in the list specified by the value of the ‘select_attr’ query parameter, in the entity body of a PUT request. [PR-12] Upon receiving such a malformed request the Provider SHALL respond with a “400 Bad Request” status code. [PR-13]

One way to think about this is to regard the values listed in ‘select_attr’ as a mask on the base resource where HTTP requests are interpreted within the context of this mask. Another way to think about this is to consider a URL that includes the ‘select_attr’ query parameter as dynamically addressing a resource that contains only those attributes whose name occurs in the value of that query parameter.

For example, the following request should be considered as an attempt to delete the ‘tags’ attribute but, because they are not listed in the value of ‘select_attr’, not an attempt to delete any other attributes:

```
PUT /my_paas/assembly/273?select_attr=tags HTTP/1.1
Host: example.org
Content-Type: application/json
Content-Length: ...
```

Note that the appearance of an attribute in the value of ‘select_attr’ that is not listed in the consumer_mutable array does not, in itself, constitute an erroneous request. As with the processing of a PUT request without ‘select_attr’, it is the **value** of the attribute in the message body that determines whether or not the Consumer is attempting to modify the resource’s attribute. A PUT request in which the value of ‘select_attr’ lists every attribute of a resource is semantically equivalent to a PUT request that does not use ‘select_attr’.

7.4.2 Updating with JSON Patch

JSON Patch [RFC6902] defines a JSON document structure for expressing a sequence of operations to apply to a JSON document, suitable for use with the HTTP PATCH method. The “application/json-patch+json” media type is used to identify such patch documents.

Providers SHALL support the HTTP PATCH method in conjunction with the “application/json-patch+json” media type with the following, additional provisions with respect to the operations defined in Section 4 of the JSON Patch specification: [PR-26]

- Providers SHALL support the ‘add’, ‘remove’, and ‘replace’ operations. [PR-27]
- Providers MAY support the ‘move’, ‘copy’, and ‘test’ operations. [PR-28]

7.4.3 Mutability of Resource Attributes

Consumers SHALL NOT send a request that attempts to change the value of a resource attribute that is not referenced by an element in the consumer_mutable array of the target resource. [PR-21] On receiving such a request the Provider SHALL generate an HTTP response with 403 HTTP status code. [PR-22]

7.5 Stopping an Application Instance

One way of stopping an application instance is to send an HTTP DELETE request to the URL of the corresponding *assembly resource*.

8 Extensions

Features provided by this specification can be extended to provide additional information and functionality. Extensions MAY be added by registering the new functionality in the *extension resource* collection. [EX-02] Extensions SHALL NOT change or remove any features or functionality of this specification. [EX-03] Each extension SHALL satisfy all the criteria in Section 9, "Conformance", and SHALL NOT contradict any normative statements in this document. [EX-04] The following extensions are allowed:

Category	Functionality	Description
API Extension	New HTTP Request Verbs	Support for additional HTTP Request Verbs that are not used by this specification, such as HEAD.
API Extension	HTTP Header Handlers	Processing of specific HTTP headers provided by clients. For example, an API Extension may require an authentication token header.
API Extension	New Resources	Addition of new resource types that MAY handle HTTP requests such as POST or PUT to create new resources of this type.
API Extension	New Resource Methods	Allow the creation of new methods or actions that may cause different sequences of state changes than happen by default.
PDP Extension	New Metadata in the PDP	Additional metadata provided in the PDP to allow for more sophisticated handling of the bundled data.
Resource Extension	New Resource Types	Addition of new resource types.
Resource Extension	New Resource Attributes	Addition of new attributes to existing resources.
Resource Extension	New States in any Application Lifecycle	Addition of new application states, such as an intermediate state between the states defined by the specification.

Table 8-1: Extension Categories and Functionality

8.1 Unique Name Requirement

Entities
<ul style="list-style-type: none">ResourcesAttributes

- Methods
- PDP Metadata Keys

Table 8-2: Entities

Entities are enumerated in Table 8-2. The Extension Developer SHALL use a unique name for new entities within an existing namespace. [EX-05] Entities added by an extension SHALL NOT interfere with names of existing entities, including any added by another extension. [EX-06]

NOTE: Each resource has its own namespace. It is acceptable to create a resource named *example.org:Foo*, and another resource named *example.org:Bar*, where both resources have an attribute named *fooBar*.

The use of your registered ICAAN Internet domain name followed by a colon (":") character as a prefix to all your entity names is RECOMMENDED to comply with these requirements. [EX-07]

Example: New Attribute "foo" added by Example Organization

```
example.org:foo
```

Example: New Attribute "foo" added by Example Inc.

```
EXAMPLE-INC:foo
```

Extension Category	New Entity	Exception
API Extension	Adding HTTP Request Verbs	Unique name not required for HTTP verbs
API Extension	Adding HTTP Header Handlers	Unique name not required for HTTP headers

Table 8-3: Unique Name Exceptions

A unique name is not required for entities listed in Table 7-3.

NOTE: RFC-3986 identifies Unreserved Characters that may be used in a URI without any encoding. Percent-Encoding allows any character to be represented in a URI. Special characters such as "." and "." have specific meanings in scripting languages such as JavaScript. Special characters must be properly escaped in order to use them as part of a name string. Your data serialization format may not escape all problematic characters, so you may need to add logic to your clients to escape special characters to enable interaction with an Extension.

8.2 extension Resource

An *extension resource* represents new functionality added to the Platform. This resource has the following pseudo schema:

```
{
  "uri": URI,
  "name": String,
  "description": String ?,
  "tags": String[] ?,
  "representation_skew": String ?,
  "external_management_resource": URI ?,
  "version": String,
  "documentation": URI ?,
  "metadata": object
}
```

The *extension resource* contains the following attributes:

8.2.1 version

Type: String

Required: true

This attribute contains a string identifier of the version of this extension.

8.2.2 documentation

Type: URI

Required: false

This attribute is a URI that references a human readable document that describes the extension.

8.3 Extending Existing Resources

New attributes MAY be added to an existing resource using an *extension resource* if the Unique Name Requirement in 8.1 is met. **[EX-10]** A new resource type is not required in order to add new attributes.

Example of an extended *extension resource*:

```
{
  "uri": URI,
  "name": String,
  "description": String,
  "version": String,
  "documentation": URI ?,
  "acme.com:foo": String ?,
  "metadata": object
}
```

9 Conformance

There are four conformance targets defined in this specification:

- CAMP Provider
- CAMP Consumer
- Platform Deployment Package
- Plan

9.1 CAMP Provider

An implementation claiming to conform to the requirements of a CAMP Provider defined in this specification SHALL comply with all of the CAMP Provider or Provider mandatory requirements listed in this specification, as summarized in Appendix B.1, “[Mandatory Statements](#)”.

A CAMP Provider supporting any CAMP Provider or Provider optional feature SHALL comply with the requirements for that feature listed in this specification and summarized in Appendix B.2 “[Non-Mandatory Statements](#)”.

9.2 CAMP Consumer

An implementation claiming to conform to the requirements of a CAMP Consumer defined in this specification SHALL comply with all of the CAMP Consumer or Consumer mandatory requirements listed in this specification, as summarized in Appendix B.1, “[Mandatory Statements](#)”.

A CAMP Consumer supporting any CAMP Consumer or Consumer optional feature shall comply with the requirements for that feature listed in this specification and summarized in Appendix B.2 “[Non-Mandatory Statements](#)”.

9.3 Platform Deployment Package

For a document to be a valid PDP, it SHALL comply with all of the PDP mandatory requirements listed in this specification, as summarized in Appendix B.1, “[Mandatory Statements](#)”.

If a PDP contains an optional PDP feature it shall comply with the requirements for that feature listed in this specification and summarized in Appendix B.2 “[Non-Mandatory Statements](#)”.

9.4 Plan

For a document to be a valid Plan, it SHALL comply with all of the Plan mandatory requirements listed in this specification, as summarized in Appendix B.1, “[Mandatory Statements](#)”.

If a Plan contains an optional Plan feature it shall comply with the requirements for that feature listed in this specification and summarized in Appendix B.2 “[Non-Mandatory Statements](#)”.

Appendix A. Acknowledgments

This appendix is non-normative. The following individuals have participated in the creation of this specification and are gratefully acknowledged:

Participants:

Roshan Agrawal	Rackspace Hosting, Inc.
Michael Behrens	US Department of Defense (DoD)
Bhaskar Reddy Byreddy	Software AG, Inc.
Mark Carlson	Oracle
Martin Chapman	Oracle
Francesco D'Andria	Cloud4SOA
Jacques Durand	Fujitsu Limited
Panagiotis Gouvas	Cloud4SOA
Keith Grange	JumpSoft
Alex Heneveld	Cloudsoft Corporation Limited
Gershon Janssen	Individual Member
David Jilk	Standing Cloud, Inc.
Duncan Johnston-Watt	Cloudsoft Corporation Limited
Anish Karmarkar	Oracle
Tobias Kunze	Red Hat
Eugene Luster	US Department of Defense (DoD)
Ashok Malhotra	Oracle
Alex McDonald	NetApp
Rich Miller	Cloudsoft Corporation Limited
Jeff Mischkinsky	Oracle
Adrian Otto	Rackspace Hosting, Inc.
Derek Palma	Vnomic
Gilbert Pilz	Oracle
Krishna Raman	Red Hat
Tom Rutt	Fujitsu Limited
Zhexuan Song	Huawei Technologies Co., Ltd.
Charles Tupitza	JumpSoft
Jeffrey West	Oracle
Prasad Yendluri	Software AG, Inc.

Appendix B. Normative Statements

B.1 Mandatory Statements

Tag	Statement
[PDP-02]	A Provider SHALL support the following archive file formats for a PDP: <ul style="list-style-type: none">• A PDP as a ZIP archive [ZIP]
[PDP-03]	A Provider SHALL support the following archive file formats for a PDP: <ul style="list-style-type: none">• A PDP as a TAR archive [TAR]
[PDP-04]	A Provider SHALL support the following archive file formats for a PDP: <ul style="list-style-type: none">• A PDP as a GZIP [RFC1952] compressed TAR archive
[PDP-10]	The format of the manifest file and the certificate file SHALL be as defined by the OVF specification [OVF].
[PDP-11]	A Platform Deployment Package (PDP) SHALL contain a single Plan file.
[PDP-27]	Providers SHALL support the “https” URI scheme as defined in RFC 2818 [RFC2818].
[PDP-29]	Providers SHALL understand this delimiter and SHALL NOT resolve any content if the archive format is unsupported.
[PLAN-01]	The Plan file SHALL be located at the root of the PDP archive.
[PLAN-02]	The Plan file SHALL be named “camp.yaml”.
[PLAN-03]	A Plan file SHALL contain a single instance of a Plan.
[PLAN-05]	For Plans that conform to this document, the value of this node SHALL be as defined in Section 1.8 “Specification Version”.
[PLAN-06]	Plans SHALL use id values that are unique within the scope of the Plan.
[PLAN-07]	Consumers SHALL follow the syntax and semantics described here when using URIs with a “pdp” scheme.
[PLAN-08]	The Plan file SHALL conform to YAML 1.1 [YAML 1.1].
[PLAN-09]	The Plan file SHALL conform to the description provided in this section.
[RE-06]	If the Required boolean constraint for an attribute of a resource type has a value of “true”, then a resource of this type SHALL have the attribute present.
[RE-07]	If a particular attribute is not listed, it indicates that the value(s) of the attribute, once set, SHALL NOT change for the lifetime of the resource.
[RE-09]	If a particular attribute is not listed, it indicates that the value(s) of the attribute SHALL NOT be changed by Consumers.
[RE-11]	If present, representation_skew SHALL have one of the following values: <ul style="list-style-type: none">• “CREATING” – describes a resource that is in the process of being created. The client can expect that the resource will have a skew of “NONE” once this process has completed.• “NONE” – is an assertion by the CAMP server that the information in the resource

	<p>is an accurate representation of the underlying platform implementation. Absent some action by the client or some other event (e.g. platform shutdown), a resource with a skew of NONE can be expected to remain in synch with the platform implementation.</p> <ul style="list-style-type: none"> • “UNKNOWN” – indicates that the CAMP server cannot accurately depict the aspect of the platform implementation represented by this resource. Users can attempt to address the underlying issues(s) by manipulating this and/or other resources as specified by the API. • “DESTROYING” – describes a resource that is in the process of being destroyed. The client can expect that the resource will cease to exist once this process has completed. 										
[RE-12]	<p>The following table lists the methods that SHALL be supported for each <code>representation_skew</code> value.</p> <table border="1"> <thead> <tr> <th><code>representation_skew</code> value</th> <th>Methods Available</th> </tr> </thead> <tbody> <tr> <td>CREATING</td> <td>GET, DELETE</td> </tr> <tr> <td>NONE</td> <td>All supported methods for that resource.</td> </tr> <tr> <td>UNKNOWN</td> <td>All supported methods for that resource.</td> </tr> <tr> <td>DESTROYING</td> <td>GET</td> </tr> </tbody> </table>	<code>representation_skew</code> value	Methods Available	CREATING	GET, DELETE	NONE	All supported methods for that resource.	UNKNOWN	All supported methods for that resource.	DESTROYING	GET
<code>representation_skew</code> value	Methods Available										
CREATING	GET, DELETE										
NONE	All supported methods for that resource.										
UNKNOWN	All supported methods for that resource.										
DESTROYING	GET										
[RE-20]	Each <i>platform_endpoint resource</i> SHALL refer to exactly one <i>platform resource</i> , and indicate the versions supported by the Platform.										
[RE-22]	For Platforms that implement this version of the CAMP specification, the value of this attribute SHALL be as defined in Section 1.8, “ Specification Version ”.										
[RE-23]	The values in this array SHALL be the Specification Version Strings of previous CAMP specification versions.										
[RE-24]	<i>platform_endpoint resources</i> that reference <i>platform resources</i> with a <code>specification_version</code> value of “CAMP 1.2” SHALL NOT include this attribute because no previous versions are compatible.										
[RE-26]	For Platforms that implement this version of the CAMP specification, the value of this attribute SHALL be as defined in Section 1.8, “ Specification Version ”.										
[RE-27]	The value of this attribute SHALL exactly match the value of the <code>specification_version</code> attribute of any <i>platform_endpoint resource</i> that references this <i>platform resource</i> .										
[RE-29]	The value of this attribute SHALL exactly match the value of the <code>implementation_version</code> attribute of any <i>platform_endpoint resource</i> that references this <i>platform resource</i> .										
[RE-37]	The collection of <i>parameter_definition resources</i> referenced by this attribute SHALL define parameters to allow setting the ‘name’, ‘description’, and ‘tags’ attributes of any new resource created in the course of interacting with this resource.										
[RE-38]	If this attribute is present in the resource, Providers SHALL support the POST method on that resource in addition to the methods defined in Section 5.5, “ HTTP Method Support ”.										
[RE-39]	An <i>assembly resource</i> SHALL have at least one <i>component resource</i> .										
[RE-40]	For every format that the Platform supports, there SHALL be a <i>format resource</i> that										

	describes that format.
[RE-41]	A <i>format resource</i> that describes the required JSON format SHALL appear in the collection of <i>format resources</i> referenced by the <i>platform resource</i> .
[RE-42]	The name, mime_type, version, and documentation attribute values for the JSON <i>format resource</i> SHALL reflect the above values.
[RE-44]	For every resource type that the Platform supports, the <i>type_definition resource</i> collection SHALL contain a <i>type_definition resource</i> that represents that a resource type.
[RE-45]	For every attribute of the type that is not inherited from its super-types, the items array SHALL contain an <i>attribute_definition resource</i> that defines that attribute.
[RE-53]	Providers SHALL support the HTTP GET method on all of the resources defined in this section.
[RE-61]	In addition to the methods defined in Section 5.5, “HTTP Method Support”, Providers SHALL support the HTTP DELETE method on the <i>assembly resource</i> .
[RE-62]	In addition to the methods defined in Section 5.5, “HTTP Method Support”, Providers SHALL support the HTTP DELETE method on the <i>component resource</i> .
[RE-64]	In addition to the methods defined in Section 5.5, “HTTP Method Support”, Providers SHALL support the HTTP POST method on the <i>operation resource</i> .
[RE-65]	Consumers and Providers SHALL express Timestamps in UTC (Coordinated Universal Time), with the special UTC designator ("Z").
[RE-70]	When supporting such a Resource, a Provider SHALL implement it and serialize it as described in the corresponding sub-section.
[RE-71]	A Consumer SHALL serialize Resource data in its requests based on the definition of this Resource as described in the corresponding sub-section.
[RE-73]	On reception of a DELETE request a Provider SHALL remove the <i>assembly resource</i> from the system along with any <i>component resources</i> referenced by that <i>assembly resource</i> . (i.e. the tree of resources that was created when the application was instantiated).
[RE-74]	On reception of a DELETE request a Provider SHALL remove the <i>assembly resource</i> from the <i>assembly_factory resource's</i> items array.
[RE-76]	In cases where a sub-type adds additional constraints to an attribute inherited from its super-types (e.g. makes an optional attribute required), Providers SHALL include an <i>attribute_definition resource</i> for that attribute in the items array.
[RE-77]	In addition to the methods defined in Section 5.5, “HTTP Method Support”, Providers SHALL support the HTTP DELETE method on the <i>plan resource</i> .
[RE-78]	On reception of a DELETE request a Provider SHALL remove the <i>plan resource</i> from the system.
[RE-79]	On reception of a DELETE request a Provider SHALL remove the <i>plan resource</i> from the <i>plan resource</i> collection.
[RE-82]	A Provider SHALL NOT list an attribute in the consumer_mutable array if it is not present in the mutable array.
[RE-83]	Providers SHALL support the HTTP PUT and PATCH methods on resources that have at least one consumer-mutable attribute.

[RE-84]	If a Provider chooses to return the ETag header in its responses, the value of this header SHALL refer to the entire collection.
[RE-85]	The value of <code>total_items</code> attribute SHALL be a non-negative integer.
[RE-86]	The value of the <code>items_per_page</code> attribute SHALL be a non-negative integer.
[RE-87]	If the <code>start_index</code> query parameter was used on the HTTP GET request, the value of this attribute SHALL equal the value of that query parameter.
[RE-88]	The value of the <code>start_index</code> attribute SHALL be a non-negative integer.
[RE-89]	The size of the array SHALL be equal to the value of the <code>items_per_page</code> attribute.
[RE-90]	If the collection is empty, the value of this attribute SHALL be an array with no elements (represented as <code>[]</code> in JSON).
[PR-01]	Providers SHALL provide representations of all available resources in JSON.
[PR-02]	Consumers and Providers SHALL NOT transmit JSON objects that contain duplicate keys.
[PR-05]	Providers SHALL respond in the requested Supported Format.
[PR-07]	If the <i>If-Match</i> header field value in the request does not match the one on the server-side, the Provider SHALL send back a '412 Precondition Failed' status code.
[PR-09]	If an attribute listed in the value of the 'select_attr' query parameter is not part of the resource, a "400 Bad Request" status code SHALL be returned.
[PR-12]	A Consumer SHALL NOT include attributes, whose name does not occur in the list specified by the value of the 'select_attr' query parameter, in the entity body of a PUT request.
[PR-13]	Upon receiving such a malformed request the Provider SHALL respond with a "400 Bad Request" status code.
[PR-18]	If a POST request body does not contain a value for a required parameter, a "400 Bad Request" response SHALL be returned.
[PR-19]	If a POST request body does not contain an acceptable value for a parameter, a "400 Bad Request" response SHALL be returned.
[PR-21]	Consumers SHALL NOT send a request that attempts to change the value of a resource attribute that is not referenced by an element in the <code>consumer_mutable</code> array of the target resource.
[PR-22]	On receiving such a request the Provider SHALL generate an HTTP response with 403 HTTP status code.
[PR-26]	Providers SHALL support the HTTP PATCH method in conjunction with the "application/json-patch+json" media type with the following, additional provisions with respect to the operations defined in Section 4 of the JSON Patch specification:
[PR-27]	Providers SHALL support the 'add', 'remove', and 'replace' operations.
[PR-29]	To support the deployment of applications using a PDP, Providers SHALL accept the media types associated with the various formats as follows: <ul style="list-style-type: none"> • ZIP: "application/x-zip"
[PR-30]	To support the deployment of applications using a PDP, Providers SHALL accept the

	media types associated with the various formats as follows: <ul style="list-style-type: none"> TAR: "application/x-tar"
[PR-31]	To support the deployment of applications using a PDP, Providers SHALL accept the media types associated with the various formats as follows: <ul style="list-style-type: none"> GZIP compressed TAR: "application/x-tgz"
[PR-32]	To support the deployment of applications using a Plan file, Providers SHALL accept the use of the "application/x-yaml" media type.
[PR-41]	TLS 1.1 [RFC4346] SHALL be implemented by the Provider.
[PR-47]	A Provider SHALL return only those attributes of the queried resource whose name occurs in the list specified by the value of 'select_attr'.
[PR-48]	On successfully processing an HTTP PUT request a Provider SHALL update all the consumer-mutable attributes of the target resource, and only these, with the values of the matching attributes in the request.
[PR-49]	Providers SHALL support the deployment of applications via HTTP POST requests on the <i>assembly_factory resource</i> as described in this section.
[PR-50]	On successfully processing the request the Provider SHALL create an <i>assembly resource</i> and return a 201 Created status code in the HTTP response.
[PR-51]	The Provider SHALL include the <i>Location</i> header in the HTTP response and the value of this header SHALL reference the newly created <i>assembly resource</i> .
[PR-52]	The Provider SHALL update the <i>items</i> attribute of the <i>assembly_factory resource</i> to include the newly created <i>assembly resource</i> .
[PR-53]	On successfully processing the request the Provider SHALL create an <i>assembly resource</i> and return a 201 Created status code in the HTTP response.
[PR-54]	The Provider SHALL include the <i>Location</i> header in the HTTP response and the value of this header SHALL reference the newly created <i>assembly resource</i> .
[PR-55]	The Provider SHALL update the <i>items</i> attribute of the <i>assembly_factory resource</i> to include the newly created <i>assembly resource</i> .
[PR-56]	Providers that support the <i>plan resource</i> and <i>plan resource</i> collection SHALL support the registration of Plans via an HTTP POST request on the <i>plan resource</i> collection as described in this section.
[PR-57]	On successfully processing the request the Provider SHALL create a <i>plan resource</i> and return a 201 Created status code in the HTTP response.
[PR-58]	The Provider SHALL include the <i>Location</i> header in the HTTP response and the value of this header SHALL reference the newly created <i>plan resource</i> .
[PR-59]	The Provider SHALL update the <i>items</i> attribute of the <i>plan resource</i> collection to include the newly created <i>plan resource</i> .
[PR-60]	Providers SHALL support the deployment of applications via HTTP POST requests on the <i>assembly_factory resource</i> in which the entity body of the request contains the PDP or Plan file that is being deployed.
[PR-61]	Providers that support <i>plan resources</i> and the <i>plan resource</i> collection SHALL support the registration of Plans via HTTP POST requests on the <i>plan resource</i> collection in which the entity body of the request contains the PDP or the Plan file that is being

	registered.
[PR-62]	On successfully processing the request the Provider SHALL create a <i>plan resource</i> and return a 201 Created status code in the HTTP response.
[PR-63]	The Provider SHALL include the <i>Location</i> header in the HTTP response and the value of this header SHALL reference the newly created <i>plan resource</i> .
[PR-64]	The Provider SHALL update the <i>items</i> attribute of the <i>plan resource</i> collection to include the newly created <i>plan resource</i> .
[PR-68]	To support the deployment of applications via a reference to either a PDP, Plan file, or <i>plan resource</i> , Providers SHALL accept the "application/json" media type.
[PR-69]	To support the registration of Plans via a reference to either a PDP or a Plan file, Providers SHALL accept the "application/json" media type.
[PR-70]	To support the registration of Plans using a PDP, Providers SHALL accept the media types associated with the various formats as follows: <ul style="list-style-type: none"> • ZIP: "application/x-zip"
[PR-71]	To support the registration of Plans using a PDP, Providers SHALL accept the media types associated with the various formats as follows: <ul style="list-style-type: none"> • TAR: "application/x-tar"
[PR-72]	To support the registration of Plans using a PDP, Providers SHALL accept the media types associated with the various formats as follows: <ul style="list-style-type: none"> • GZIP compressed TAR: "application/x-tgz"
[PR-73]	To support the registration of Plans using a Plan file, Providers SHALL accept the use of the "application/x-yaml" media type.
[PR-74]	Providers SHALL support the deployment of applications via HTTP POST requests on the <i>assembly_factory resource</i> as described in this section.
[PR-75]	Providers that support <i>plan resources</i> and the <i>plan resource</i> collection SHALL support the registration of Plans via HTTP POST requests on the <i>plan resource</i> collection as described in this section.
[PR-76]	To allow an update of a subset of a resource's attributes, Providers SHALL support the use of the 'select_attr' query parameter in conjunction with the HTTP PUT method.
[PR-78]	A Provider SHALL return only those attributes of the resources in the collection whose name occurs in the list specified by the value of 'select_collection_attr'.
[PR-79]	If an attribute listed in the value of the 'select_collection_attr' query parameter is not part of the resource, it SHALL not appear in the returned resource.
[PR-80]	If a resource in the collection contains none of the attributes listed in the 'select_collection_attr' query parameter, then it SHALL be returned as an empty object.
[PR-82]	If the 'select_collection_attr' query parameter is supplied for a URL that is not a collection resource, the provider SHALL return 400 Bad Request.
[PR-83]	Providers SHALL remove all duplicates from <i>collection resources</i> and return a unique collection of sets of attributes.
[PR-84]	In cases where pagination is used in combination with 'select_collection_attr', the Provider SHALL apply pagination to the unique collection of sets of attributes.

[EX-03]	Extensions SHALL NOT change or remove any features or functionality of this specification.
[EX-04]	Each extension SHALL satisfy all the criteria in Section 9, “ Conformance ”, and SHALL NOT contradict any normative statements in this document.
[EX-05]	The Extension Developer SHALL use a unique name for new entities within an existing namespace.
[EX-06]	Entities added by an extension SHALL NOT interfere with names of existing entities, including any added by another extension.
[EX-08]	For every available extension, the <i>extension resource</i> collection referenced by the <i>platform resource</i> SHALL contain an <i>extension resource</i> that represents the extension.
[RMR-01]	If a Consumer includes this node in a Plan, the value of this node SHALL reference a Consumer-visible resource within the target Platform.
[RMR-02]	Providers SHALL support the HTTP POST method on the <i>assembly_factory resource</i> as described in Section 7.1, “ Deploying an Application ”.
[RMR-03]	The <i>assembly_factory resource</i> SHALL reference a collection of <i>parameter_definition resources</i> that describe the <i>pdp_uri</i> , <i>plan_uri</i> , <i>pdp_file</i> , <i>plan_file</i> , <i>name</i> , <i>description</i> , and <i>tags</i> parameters.
[RMR-04]	Providers that support Plans SHALL include this attribute in all <i>assembly resources</i> .
[RMR-05]	In addition to the methods defined in Section 5.5, “ HTTP Method Support ”, Providers SHALL support the HTTP POST method on the <i>plan_factory resource</i> as described in Section 7.2, “ Registering a Plan ”.
[RMR-06]	The <i>plan_factory resource</i> SHALL reference a collection of <i>parameter_definition resources</i> that describe the <i>pdp_uri</i> , <i>plan_uri</i> , <i>pdp_file</i> , <i>plan_file</i> , <i>name</i> , <i>description</i> , and <i>tags</i> parameters.
[RMR-07]	The schema of the <i>plan resource</i> returned from a CAMP Provider SHALL conform to the schema for Plans described in Section 4.3, “ Plan Schema ”, with the following additional requirements:
[RMR-08]	Representations of the <i>plan resource</i> SHALL be serialized as JSON, unless another format is negotiated.
[RMR-11]	Regardless of whether a Consumer attempts to create an <i>assembly resource</i> by POSTing to the <i>assembly_factory resource</i> or creates a <i>plan resource</i> by POSTing to the <i>plan_factory resource</i> , a Provider that supports the <i>plan resources</i> and the <i>plan_factory resource</i> SHALL create a <i>plan resource</i> for every deployed application.
[RMR-12]	<p>Providers that support <i>plan resources</i> and the <i>plan resource</i> collection SHALL advertise such support using the following <i>extension resource</i>: [RMR-12]</p> <pre> { "uri": <as appropriate>, "name": "CAMP Plans Extension", "type": "extension", "description": "indicates support for plans and plan resources", "version": "CAMP 1.2", "documentation": "http://docs.oasis-open.org/camp/camp- spec/v1.2/camp-spec-v1.2.pdf" } </pre>

[RMR-13]	Providers SHALL support PDPs that use either the ZIP [ZIP], TAR [TAR], or GZIP [RFC1952] compressed TAR formats.
[MO-02]	A sub-type SHALL NOT loosen the constraints of an attribute inherited from its super-type(s).
[MO-03]	A resource type MAY inherit from more than one super-type.
[MO-04]	If there is an attribute name collision when a sub-type inherits from multiple super-types, the inherited attributes of the same name SHALL NOT contradict the constraints and semantics of the attributes defined in its super-types.
[MO-05]	All CAMP resources SHALL inherit directly or indirectly from this resource.
[MO-06]	Items in this collection SHALL NOT, either directly or transitively, refer to the described resource.
[OP-02]	The Provider SHALL use the value of the 'sort' query parameter to order the elements of the <code>items</code> array in the returned <i>collection resource</i> .
[OP-03]	In cases where pagination is used in combination with sorting, the Provider SHALL perform the sort operation before applying pagination.
[OP-04]	A Provider SHALL respond to a request in which the value of the 'sort' query parameter specifies one or more non-scalar attributes (e.g. an attribute whose value is an array) with a "400 Bad Request" status code.
[OP-06]	If provided, the value of the 'start_index' query parameter SHALL be a non-negative integer.
[OP-07]	If provided, the value of the 'max_page' query parameter SHALL be a positive integer.
[OP-09]	A Provider SHALL NOT respond to an HTTP GET request on a collection resource with a representation whose <code>items_per_page</code> value is greater than the value of the 'max_page' query parameter.
[OP-10]	A Provider SHALL respond to a request in which the value of the 'start_index' query parameter is equal to or greater than the number of items in the collection with a "400 Bad Request" status code.
[OP-12]	If the specified resource is not present in the collection a "404 Not Found" status code SHALL be returned by the Provider.
[OP-13]	The <code>items_per_page</code> attribute SHALL be 1.
[OP-14]	The <code>items</code> array SHALL contain the resource referenced by value of the query parameter.

B.2 Non-Mandatory Statements

Tag	Statement
[PDP-01]	A PDP archive MAY include other files related to the application including, but not limited to, language-specific bundles, resource files, application content files such as web archives, database schemas, scripts, source code, localization bundles, and icons; and metadata files such as manifests, checksums, signatures, and certificates.
[PDP-05]	Providers MAY support additional archive file formats for the PDP.
[PDP-06]	A PDP MAY contain a manifest file, named <code>camp.mf</code> , at the root of the archive.

[PDP-07]	A Provider SHOULD reject a PDP if any digest listed in the manifest does not match the computed digest for that file in the package.
[PDP-08]	A PDP MAY contain a certificate, named <code>camp.cert</code> , at the root of the archive.
[PDP-09]	A Provider SHOULD reject any PDP for which the signature verification fails.
[PDP-14]	Providers MAY reflect the value of this attribute in the names of any resources that are created in the processing the <i>plan</i> .
[PDP-15]	Providers MAY reflect the value of this attribute in the descriptions of the resources that are in the processing the <i>plan</i> .
[PDP-16]	Providers MAY reflect the values of this attribute in the tags of the resources that are created in the processing of the <i>plan</i> .
[PDP-30]	Plans, Artifact Specifications, and Service Specifications MAY contain the following nodes:
[PDP-22]	The artifact MAY be contained within the PDP or MAY exist in some other location.
[PDP-28]	A Provider MAY support additional URI schemes listed at http://www.iana.org/assignments/uri-schemes/uri-schemes.xhtml .
[RE-08]	If a particular attribute is listed, it indicates that the value of the attribute MAY change due to the actions or activity of either the Provider or the Consumer.
[RE-10]	If a particular attribute is listed, it indicates that Consumers MAY change the value of the attribute.
[RE-13]	For each <code>representation_skew</code> value, CAMP Providers MAY support HTTP methods in addition to those listed in the corresponding row of Table 5-1.
[RE-15]	A Provider MAY concurrently offer multiple instances of the CAMP API.
[RE-17]	A Provider MAY expose the <i>platform_endpoints_resource</i> collection and corresponding <i>platform_endpoint_resources</i> in a way that allows for version discovery before the client has authenticated.
[RE-25]	Multiple implementations of the same CAMP specification MAY be offered concurrently.
[RE-28]	A Provider MAY choose to offer multiple implementations of the same CAMP specification.
[RE-47]	The Operation MAY require content in the body of the POST, such as parameters.
[RE-48]	The response to a POST request on an <i>operation_resource</i> SHOULD indicate what changes were made on the target resource.
[RE-49]	For asynchronous operations, the response SHOULD indicate how to track the progress of the request operation.
[RE-51]	When a “value” attribute is supplied, any timestamp provided in this attribute SHOULD correspond to when that value was observed.
[RE-52]	Extensions MAY be defined to govern common sensor management operations, such as enabling, disabling, adjusting collection frequency, specifying the history of values which should be remembered, or collecting immediately.
[RE-54]	Providers MAY elect to support additional HTTP methods in addition to those described here.
[RE-68]	This attribute MAY have one of the following values: <ul style="list-style-type: none"> • “RUNNING” – indicates that the component is functioning as expected.

	<ul style="list-style-type: none"> • “ERROR” – indicates that the component has encountered some sort of error.
[RE-69]	Providers MAY extend this list with additional values.
[RE-80]	A Provider MUST only return representations of <i>collection resources</i> whose <i>items</i> array contains resources whose type is either that described by the <i>type_definition resource</i> referenced by the <i>collection_type</i> attribute or a type derived, directly or indirectly, from that type.
[RE-81]	Metadata for a resource is extensible; a Platform Provider MAY provide additional metadata attributes.
[PR-03]	If a Consumer sends a Provider a request containing duplicate keys in a JSON object, the Provider SHOULD reject the request by sending back a ‘400 Bad Request’ status code.
[PR-04]	If a Provider sends a Consumer a response containing duplicate keys in a JSON object, the Consumer SHOULD raise an error to the user indicating the response from the server was malformed.
[PR-06]	All PUT requests that update a resource SHOULD contain the <i>If-Match</i> header field with a single entity tag value.
[PR-08]	To retrieve a subset of the attributes in a resource, the Consumer MAY use the ‘select_attr’ query parameter in conjunction with the HTTP GET method.
[PR-10]	The “select_attr” query parameter MAY appear more than once (separated by an “&”).
[PR-14]	Parameters MAY be included when performing a POST request on any resource with a <i>parameter_definitions_collection</i> attribute defined.
[PR-15]	Parameters MAY have the same name as an attribute of the resource.
[PR-16]	In such cases the Provider SHOULD set the value of that attribute to the value of the parameter OR clearly document alternate behavior.
[PR-17]	The <i>parameter_extension</i> attribute MAY be used to reference the extension which documents how the parameter is handled.
[PR-20]	All HTTP responses that return representation of a resource SHOULD use a strong ETag response header field indicating the current value of the entity tag for the resource.
[PR-23]	Consumers MAY use the HTTP PUT method to replace the representation of a resource, in its entirety, with a new representation that adds, omits or replaces the values for some of the attributes.
[PR-24]	Alternatively, Consumers MAY use the HTTP PATCH [RFC5789] method and the “application/json-patch+json” media type [RFC6902] to add, delete, or replace specific attributes.
[PR-25]	If a resource attribute is present on a resource and if an HTTP PUT request omits that attribute, it SHOULD be treated by the Provider as a request to delete the attribute.
[PR-28]	Providers MAY support the ‘move’, ‘copy, and ‘test’ operations.
[PR-33]	The JSON object MAY contain additional name-value pairs that are not defined in this specification.
[PR-40]	Requests sent from Consumers across unsecured networks SHOULD use the HTTPS protocol.
[PR-42]	TLS 1.2 [RFC5246] is RECOMMENDED.

[PR-43]	When TLS is implemented, the following cipher suites are RECOMMENDED to ensure a minimum level of security and interoperability between implementations: <ul style="list-style-type: none"> • TLS_RSA_WITH_AES_128_CBC_SHA (mandatory for TLS 1.1/1.2)
[PR-44]	When TLS is implemented, the following cipher suites are RECOMMENDED to ensure a minimum level of security and interoperability between implementations: <ul style="list-style-type: none"> • TLS_RSA_WITH_AES_256_CBC_SHA256 (addresses 112-bit security strength requirements)
[PR-45]	For each Supported Format, Consumers MAY request any resource from the Provider in that format.
[PR-46]	The JSON object MAY contain additional name-value pairs that are not defined in this specification.
[PR-77]	To retrieve a subset of the attributes of the resources in a collection, the Consumer MAY use the 'select_collection_attr' query parameter in conjunction with the HTTP GET method.
[PR-81]	The 'select_collection_attr' query parameter MAY appear more than once (separated by an "&"), in which case the value of the query parameter is considered to be the union of the attribute sets specified by each individual query parameter.
[EX-02]	Extensions MAY be added by registering the new functionality in the <i>extension resource</i> collection.
[EX-07]	The use of your registered ICAAN Internet domain name followed by a colon (":") character as a prefix to all your entity names is RECOMMENDED to comply with these requirements.
[EX-10]	New attributes MAY be added to an existing resource using an Extension if the Unique Name Requirement in 8.1 is met.
[RMR-09]	Any href attributes of ServiceSpecifications SHOULD refer to a Service resource.
[RMR-10]	All href attributes in the <i>plan resource</i> SHOULD be set to a consumer accessible URL. If the original Plan file referred to a local file, the URL indicates where the Provider stored the content.
[MO-01]	A sub-type MAY further restrict the constraints of an attribute inherited from its super-type(s).
[MO-07]	If a type inherits only from the <i>camp_resource</i> type then this attribute MAY be absent.
[OP-01]	The Consumer MAY specify the 'sort' query parameter when sending an HTTP GET request to a <i>collection resource</i> .
[OP-05]	A Consumer MAY specify the 'start_index' and 'max_page' query parameters when sending an HTTP GET request to a <i>collection resource</i> .
[OP-08]	A Provider MAY respond to an HTTP GET request on a <i>collection resource</i> with a representation whose <i>items_per_page</i> value is less than the number of items requested by the Consumer (whether explicitly via the 'max_page' query parameter or implicitly through the absence of this query parameter).
[OP-11]	To retrieve the index of an item within a collection, the Consumer MAY use the <i>index_in_collection</i> query parameter in conjunction with the HTTP GET method.

Appendix C. Example Version Scheme

This appendix is non-normative. The table below describes the version semantics from the Maven POM Syntax [[POM-Syntax](#)] for your consideration. In the absence of a prevailing version range scheme for your types, this approach is suggested. Range values can be used as version string values to convey matching semantics to the platform. This is useful when using a `CharacteristicSpecification` (see Section 4.3.7 “[CharacteristicSpecification](#)”) that can be satisfied by more than one potential match.

Range	Meaning
1.0	$x \geq 1.0$
(,1.0]	$x \leq 1.0$
(,1.0)	$x < 1.0$
[1.0]	$x == 1.0$
[1.0,)	$x \geq 1.0$
(1.0,)	$x > 1.0$
(1.0,2.0)	$1.0 < x < 2.0$
[1.0,2.0]	$1.0 \leq x \leq 2.0$
(,1.0],[1.2,)	$x \leq 1.0$ or $x \geq 1.2$. Multiple sets are comma-separated.
(,1.1),(1.1,)	$x \neq 1.1$

Table D-1: POM Version Semantics

In cases where a prevailing scheme is already popular for a given technology, use the prevailing scheme.

Appendix D. Revision History

Revision	Date	Editor	Changes Made
50	2015-04-02	Gilbert Pilz	Applied resolution to CAMP-176 . Changed version number to "x.x" pending resolution of CAMP-182 .
01	2015-04-08	Gilbert Pilz	Changed version number to "1.2" to resolve CAMP-182 . Restarting working draft numbering to coincide with this change. Resolved CAMP-178 by requiring Providers to support the HTTP DELETE method on plan resources. Miscellaneous editorial changes.
02	2015-04-22	Gilbert Pilz	Resolved CAMP-179 . Miscellaneous editorial changes.
03	2015-04-30	Gilbert Pilz	Resolved CAMP-181 .
04	2015-05-27	Gilbert Pilz	Resolved CAMP-177 ; updated diagrams.
05	2015-07-08	Gilbert Pilz, Tom Rutt	Resolved CAMP-192 ; updated diagrams. Miscellaneous editorial changes.
06	2015-08-05	Gilbert Pilz	Resolved CAMP-188 . Miscellaneous editorial changes.
07	2016-03-02	Gilbert Pilz	Resolved CAMP-180 and CAMP-201 . Miscellaneous editorial changes.
08	2016-03-31	Gilbert Pilz	Resolved CAMP-202 . Cleaned up some requirements and text that still used to the pre-07 style of defining mutability and consumer-mutability.
09	2016-05-05	Gilbert Pilz	Resolved CAMP-205 . Cleaned up some lingering references (including requirement RE-75) that referred to the now-defunct 'type' attribute.
10	2016-05-19	Gilbert Pilz	Resolved CAMP-195 . Miscellaneous editorial changes.
11	2016-05-26	Gilbert Pilz	Resolved CAMP-210 . Miscellaneous editorial changes.
12	2016-09-29	Gilbert Pilz	Resolved CAMP-186 and CAMP-187 . Miscellaneous editorial changes.
13	2016-11-16	Gilbert Pilz	Resolved CAMP-193 and CAMP-204 . Miscellaneous editorial changes.
14	2017-02-16	Gilbert Pilz	Resolved CAMP-212 , CAMP-215 , CAMP-216 , and CAMP-219 . Miscellaneous editorial changes.
15	2017-03-01	Gilbert Pilz	Resolve CAMP-198 . Miscellaneous editorial

			changes.
16	2017-03-10	Gilbert Pilz	Fix Figure 5-1 to reflect the fact that <i>platform_endpoints</i> is a collection resource. Miscellaneous editorial changes and cleanups.
17	2017-05-05	Gilbert Pilz	Resolved CAMP-221 . Updated UML diagrams to reflect changes to the resource model from CAMP 1.1. Miscellaneous editorial changes and cleanups.
18	2017-09-21	Gilbert Pilz	Resolved CAMP-222 , CAMP-225 , CAMP-227 , CAMP-228 , CAMP-229 , CAMP-230 , CAMP-231 , CAMP-232 , CAMP-233 , CAMP-234 , CAMP-236 , CAMP-237 , CAMP-238 , CAMP-239 , and CAMP-40 . Miscellaneous editorial changes and cleanups.