



1

2

3

Architecture for Managing Clouds

4

A White Paper from the Open Cloud Standards Incubator

5

Version: 1.0.0

6

Status: DMTF Informational

7

Publication Date: 2010-06-18

8

Document Number: DSP-IS0102

9 Copyright © 2010 Distributed Management Task Force, Inc. (DMTF). All rights reserved.

10 DMTF is a not-for-profit association of industry members dedicated to promoting enterprise and systems
11 management and interoperability. Members and non-members may reproduce DMTF specifications and
12 documents, provided that correct attribution is given. As DMTF specifications may be revised from time to
13 time, the particular version and release date should always be noted.

14 Implementation of certain elements of this standard or proposed standard may be subject to third party
15 patent rights, including provisional patent rights (herein "patent rights"). DMTF makes no representations
16 to users of the standard as to the existence of such rights, and is not responsible to recognize, disclose,
17 or identify any or all such third party patent right, owners or claimants, nor for any incomplete or
18 inaccurate identification or disclosure of such rights, owners or claimants. DMTF shall have no liability to
19 any party, in any manner or circumstance, under any legal theory whatsoever, for failure to recognize,
20 disclose, or identify any such third party patent rights, or for such party's reliance on the standard or
21 incorporation thereof in its product, protocols or testing procedures. DMTF shall have no liability to any
22 party implementing such standard, whether such implementation is foreseeable or not, nor to any patent
23 owner or claimant, and shall have no liability or responsibility for costs or losses incurred if a standard is
24 withdrawn or modified after publication, and shall be indemnified and held harmless by any party
25 implementing the standard from any and all claims of infringement by a patent owner for such
26 implementations.

27 For information about patents held by third-parties which have notified the DMTF that, in their opinion,
28 such patent may relate to or impact implementations of DMTF standards, visit
29 <http://www.dmtf.org/about/policies/disclosures.php>.

30

31 **Abstract**

32 This white paper is one of two Phase 2 deliverables from the DMTF Cloud Incubator and describes the
33 reference architecture as it relates to the interfaces between a cloud service provider and a cloud service
34 consumer. The goal of the Incubator is to define a set of architectural semantics that unify the
35 interoperable management of enterprise and cloud computing.

36 **Acknowledgments**

37 The Open Cloud Standards Incubator wishes to acknowledge contributions from the following people:

- 38 • Kane, Valerie – Advanced Micro Devices
- 39 • Vancil, Paul – Advanced Micro Devices
- 40 • Kowalski, Vincent – BMC Software
- 41 • Lipton, Paul – CA, Inc.
- 42 • Moscovich, Efraim – CA, Inc.
- 43 • Waschke, Marvin – CA, Inc.
- 44 • Armstrong, Joseph – Cisco
- 45 • Sankar, Krishna – Cisco
- 46 • Serr, Frederick – Cisco
- 47 • Siefer, Mike – Cisco
- 48 • Tomsu, Peter – Cisco
- 49 • Wheeler, Jeff – Cisco
- 50 • Pardikar, Shishir – Citrix Systems Inc.
- 51 • Sirjani, Abolfazl – Citrix Systems Inc.
- 52 • Landau, Richard – Dell
- 53 • Ericson, George – EMC
- 54 • Linn, John – EMC
- 55 • Durand, Jacques – Fujitsu
- 56 • Song, Zhexuan – Fujitsu
- 57 • Aguiar, Glaucimar – Hewlett-Packard Company
- 58 • Cook, Nigel – Hewlett-Packard Company
- 59 • Maciel, Fred – Hitachi, Ltd.
- 60 • Baskey, Michael – IBM
- 61 • Johnson, Mark – IBM
- 62 • Pendarakis, Dimitrios – IBM
- 63 • Cox, Billy – Intel Corporation
- 64 • Li, Hong – Intel Corporation

- 65 • Raghu, Rekha – Intel Corporation
- 66 • Parchem, John – Microsoft Corporation
- 67 • Carter, Steve – Novell
- 68 • Jose, Jaimon – Novell
- 69 • Carlson, Mark – Oracle
- 70 • Vambenepe, William – Oracle
- 71 • Yu, Jack – Oracle
- 72 • Owens, Ken – Savvis
- 73 • Qualls, Richard – Sun Microsystems, Inc.
- 74 • Norbeck, Don – SunGard Availability Services
- 75 • Ronco, Enrico – Telecom Italia
- 76 • Galán Márquez, Fermín – Telefónica
- 77 • Bumpus, Winston – VMware Inc.
- 78 • Lamers, Lawrence – VMware Inc.
- 79 • Matthews, Jeanna – VMware Inc.
- 80
- 81

CONTENTS

83	1	Scope	7
84	2	References	7
85	3	Terms and Definitions	8
86	4	Symbols and Abbreviated Terms.....	12
87	5	Introduction.....	14
88	6	Reference Architecture Overview	16
89	6.1	Service Lifecycle and Use Cases	16
90	6.2	Resource Models	17
91	6.3	Interaction Patterns.....	18
92	6.4	Security Architecture.....	20
93	6.5	Role of Rules, Policies, and Constraints.....	21
94	7	Requirements	23
95	7.1	Protocol Stack.....	23
96	7.2	Resource Model.....	23
97	7.3	Adoptability	24
98	7.4	Internationalization.....	24
99	7.5	Rules, Constraints, and Policies	24
100	7.6	Cloud Management Interface Security	25
101	7.7	Data and Storage.....	26
102	7.8	Logs, Event Management, Incident Response, and Notification	26
103	7.9	Audit, Legal, and Compliance Monitoring	27
104	7.10	Security Considerations for Virtualization Technology	27
105	7.11	Portability and Interoperability for Secure Migration	27
106	8	Protocol Examples	28
107	8.1	Message Exchange Patterns	28
108	8.2	Infrastructural Aspects Affecting the Choice of MEP and Its Execution Mode	34
109	9	Cloud Ecosystem	34
110	9.1	Architectural Impact on Data Artifacts	36
111	9.2	Cloud Service Provider Identity/Key Store [200]	37
112	9.3	Cloud Service Provider Policy Store [205].....	38
113	9.4	Cloud Service Provider Service Store [210]	39
114	9.5	Cloud Service Provider Billing/Compliance Event Store [215]	39
115	9.6	Cloud Service Provider Deployment [300].....	40
116	10	Conclusion and Next Steps	40
117	10.1	Standards Development Steps	40
118	10.2	Important Extensions to Be Considered	40
119	10.3	Integration with Alliance Partner Frameworks	41
120	11	Future of the Cloud Incubator	41
121	Annex A (informative)	Message Exchange Protocol Examples.....	42
122	Annex B (informative)	Policy Discussion	52
123	Annex C (informative)	Change Log.....	57

125 **Figures**

126 Figure 1 – Scope and Benefits of DMTF Open Cloud Standards Incubator..... 15

127 Figure 2 – Cloud Service Reference Architecture 16

128 Figure 3 – Cloud Service Lifecycle States and Use Cases 17

129 Figure 4 – Interaction Patterns..... 19

130 Figure 5 – Example of Interaction Patterns for a Provision Service Use Case 20

131 Figure 6 – Security Context..... 20

132 Figure 7 – Constraints Flow 22

133 Figure 8 – One-Way MEP Execution Modes 30

134 Figure 9 – Two-Way/Sync MEP Execution Mode 31

135 Figure 10 – Two-Way/Push-and-Push MEP Execution Mode 32

136 Figure 11 – Two-Way/Push-and-Pull MEP Execution Mode 33

137 Figure 12 – High-Level Architecture 35

138 Figure 13 – Expanded Architecture 36

139 Figure B-1 – Policy Model 52

140

141 **Tables**

142 Table 1 – Cross-Mapping of Elements in Figure B.1 with Elements in Figure 13 38

143 Table B-1 – Policy Examples 54

144

145 1 Scope

146 This document is one of two documents that together describe how standardized interfaces and data
147 formats can be used to manage clouds. This document focuses on the overall architecture; the other
148 document focuses on interactions and data formats.

149 The scope of this architecture document includes:

- 150 • **Reference architecture for managing clouds.** This reference architecture was introduced in
151 the DMTF *Interoperable Clouds* white paper ([DSP-IS0101](#)). The concepts are further explored
152 and described in this document.
- 153 • **Requirements.** These requirements are for the architected interfaces in general, including
154 requirements for the protocols, resource model, and security mechanisms.
- 155 • **Role of policies and constraints.** Given the focus on the interfaces for managing clouds rather
156 than on the internal details of a cloud implementation, a useful abstraction is to define the
157 desired capabilities of the cloud using policies and constraints. These are interpreted to be *what*
158 the cloud service provider should offer rather than *how* it offers it.
- 159 • **Patterns for interactions.** These patterns of consumer/provider interactions repeat across
160 many use cases with different operations and data payloads, depending on the use case.
- 161 • **An example cloud ecosystem.** To aid comprehension, an example of how a cloud service
162 provider might design the management interfaces is provided, with a discussion of design
163 considerations. It should be emphasized that this section is not prescriptive; rather, it is
164 illustrative.

165 A companion white paper to this document, *Use Cases and Interactions for Managing Clouds* ([DSP-](#)
166 [IS0103](#)), describes other aspects of managing clouds, namely:

- 167 • management use cases across the entire lifecycle of a cloud service
- 168 • interaction sequences among consumers, developers, and providers to implement the use
169 cases
- 170 • data artifacts exchanged in the interaction sequences

171 The following aspects of clouds are specifically out of the scope of this document:

- 172 • The architecture addresses management function only; it does not address how general
173 business applications use business services exposed by a cloud, or the applications deployed
174 into a cloud.
- 175 • The architecture does not address how to *build* management function in a cloud. Instead, it
176 focuses on the management *interfaces* to the cloud. Providers are free to implement the
177 services behind these interfaces in any way.

178 2 References

179 Cloud Security Alliance, *Security Guidance for Critical Areas of Focus in Cloud Computing 2.1*,
180 <http://www.cloudsecurityalliance.org/csaguide.pdf>

181 DMTF DSP0243, *Open Virtualization Format Specification 1.1*,
182 http://www.dmtf.org/standards/published_documents/DSP0243_1.1.pdf

183 DMTF DSP1001, *Management Profile Specification Usage Guide 1.0*,
184 http://www.dmtf.org/standards/published_documents/DSP1001_1.0.pdf

- 185 DMTF DSP-IS0101, *Interoperable Clouds, A White Paper from the Open Cloud Standards Incubator 1.0*,
186 http://www.dmtf.org/standards/published_documents/DSP-IS0101_1.0.pdf
- 187 DMTF DSP-IS0103, *Use Cases and Interactions for Managing Clouds, A White Paper from the Open
188 Cloud Standards Incubator 1.0*, http://www.dmtf.org/standards/published_documents/DSP-IS0103_1.0.pdf
- 189 ITIL® V3 Glossary, *Glossary of Terms, Definitions, and Acronyms*, v3.1.24, 11 May 2007,
190 http://www.best-management-practice.com/gempdf/ITIL_Glossary_V3_1_24.pdf
- 191 National Institute of Standards and Technology, *The NIST Definition of Cloud Computing*,
192 <http://csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.doc>
- 193 OASIS, *ebXML Messaging Services Version 3.0*,
194 http://www.oasis-open.org/committees/download.php/24618/ebms_core-3.0-spec-cs-02.pdf
- 195 OASIS, *Web Services Make Connection (WS-MakeConnection) Version 1.1*, <http://docs.oasis-open.org/ws-rx/wsmc/200702/wsmc-1.1-spec-os.html>
- 196
- 197 Storage Networking Industry Association, *Cloud Data Management Interface Version 1.0g*,
198 http://www.snia.org/tech_activities/publicreview/Cloud_Data_Management_Interface_Draft1.0.pdf

199 3 Terms and Definitions

200 For the purposes of this document, the following terms and definitions apply.

201 3.1

202 **client**

203 runs proprietary cloud software stacks side-by-side with DMTF standards-based protocol software stacks

204 3.2

205 **cloud service**

206 a publicly available service or a private service that is used within an enterprise

207 3.3

208 **cloud service consumer**

- 209 • approves business/financial expenditures for consumed services
- 210 • maintains accounts for used service instances
- 211 • requests service instances and changes to service instances (typically on behalf of the consumer
212 business manager)
- 213 • provides access to services for service users

214 See Figure 2.

215 3.4

216 **cloud service developer**

217 designs, implements, and maintains service templates (see Figure 2)

218 3.5

219 **cloud service provider**

220 an organization that supplies cloud services to one or more internal or external consumers (see Figure 2)

221 3.6

222 **configuration management database**

223 **CMDB**

224 as used in the IT Infrastructure Library® context, a repository of information that represents authorized
225 configuration of all the components of an information system

- 226 **3.7**
227 **data artifacts**
228 as used in this document, the control and status elements exchanged across the [provider interface](#) using
229 the infrastructure
- 230 **3.8**
231 **deployment**
232 the process of creating a service instance in a reserved or prepared environment
233 NOTE: This may be more suitable for software or configuration deployment to existing running services.
- 234 **3.9**
235 **infrastructure**
236 as used in this document, the means by which control and status elements are exchanged between the
237 [cloud service provider](#) and the [cloud service consumer](#)
238 The elements so exchanged are referred to as [data artifacts](#).
- 239 **3.10**
240 **Infrastructure as a Service**
241 **IaaS**
242 the capability provided to the consumer to provision processing, storage, networks, and other
243 fundamental computing resources where the consumer is able to deploy and run arbitrary software, which
244 can include operating systems and applications
245 The consumer does not manage or control the underlying cloud infrastructure but has control over
246 operating systems, storage, deployed applications, and possibly limited control of select networking
247 components (for example, host firewalls). (Source: [NIST](#))
- 248 **3.11**
249 **notification**
250 a signal from the [cloud service provider](#) to the [cloud service consumer](#) that a condition exists that requires
251 attention
252 A notification can be either polled or sent asynchronously.
- 253 **3.12**
254 **profiles**
255 a specification that defines the CIM model and associated behavior for a management domain (for
256 example, Server Virtualization) (Source: [DSP1001](#))
257 The CIM model includes the CIM classes, associations, indications, methods, and properties. The
258 management domain is a set of related management tasks.
- 259 **3.13**
260 **provider interface**
261 the interface through which [cloud service consumers](#) may access and monitor their contracted services
262 The interface covers [SLO](#) negotiation and measurement, service access, service monitoring, and billing.
263 This interface is also the interface through which a [cloud service developer](#) interacts with a [cloud service](#)
264 [provider](#) to create a service template that could be added to the service catalog. Service offerings, which
265 contain the service template, are configured by providers. Cloud service consumers can then select the
266 service offerings for their use. See Figure 2.

267 3.14**268 provisioning**

269 the process of selecting, reserving, or creating an instance of a service offering

270 Service offerings are selected from the service provider's service catalog and are then provisioned into
271 service instances.

272 Provisioning is also the process of selecting or reserving service resources from available pools,
273 assembling them together, and configuring them based on a specific request in the contract, in order to
274 fulfill the contract.

275 For example, a server instance can be created from a template; assigned CPU, memory, storage, and
276 network resources; and configured for the consumer to satisfy the contract requirements (apply patches,
277 adjust security, configure firewall, and so on).

278 3.15**279 security manager**

280 responsible for managing the credentials and authentication processes as they relate to the operations
281 across the [provider interface](#)

282 Security requirements for the cloud include user authentication, identity and access management, data
283 protection, multi-tenancy resource isolation, monitoring and auditing for compliance, incident response,
284 user and customer privacy, and the underlying portability and interoperability of security components.

285 3.16**286 service catalog**

287 a database of information about the cloud services offered by a service provider

288 The service catalog includes a variety of information about the services, including description of the
289 services, the types of services, cost, supported [SLAs](#), and who can view or use the services. More
290 generally, the service catalog contains information about services through their entire lifecycle. It contains
291 service templates (created by developers), service offerings (created by providers), and deployed service
292 instances.

293 3.17**294 service contract**

295 an agreement between the [cloud service provider](#) and [cloud service consumer](#) to state the terms of
296 service usage by the cloud service consumer

297 3.18**298 service entity**

299 the representation of an identifiable logical element that serves and satisfies a list of operations

300 For example, a server (both physical and virtual) that contains an OS stack is a service entity that
301 supports the prescribed operations defined on the particular OS. In addition, a service entity can contain a
302 composite of other service entities or elements, as long as the group of these entities and elements
303 satisfies, collectively, the prescribed operations.

304 3.19**305 service instance**

306 the instantiation of a service request

307 3.20**308 service offering**

309 a service template combined with service-level agreements, constraints, costs, billing information and
310 other data necessary to offer the service described in the template to a consumer

- 311 **3.21**
312 **service request**
313 a request by a consumer to instantiate a service offering. Requests require service contracts, which may
314 be created prior to or simultaneously with service requests
- 315 **3.22**
316 **service template**
317 a collection of items (machine images, connectivity definitions, storage, and so on) that are stored in the
318 [service catalog](#) and can be provisioned at the [cloud service provider](#) (see Figure 3)
- 319 **3.23**
320 **service-level agreement**
321 **SLA**
322 in the context of [cloud service providers](#), a negotiated, legally binding contract between the cloud service
323 provider and [cloud service consumer](#)
- 324 The SLA includes agreements about services and responsibilities for service availability, performance,
325 and billing. The SLA must be structured such that it can be propagated to all involved cloud service
326 providers. The consumer contracts services from one cloud service provider who could be a broker,
327 federator, or non-broker/non-federator, and that SLA is used as the basis for the broker or federator to
328 programmatically contract for services from other cloud service providers.
- 329 **3.24**
330 **service-level objective**
331 **SLO**
332 a unique, distinct, and measurable aspect of an [SLA](#)
- 333 All parties of the SLA must agree that sets of SLOs and their measurement, correlation, and reporting will
334 represent the compliance or non-compliance of unique stipulations within the SLA. The output of one or
335 more correlated SLOs must be deterministic and supported by audit logs that allow interested parties to
336 prove the output of an SLO or the correlation of the output of several SLOs.
- 337 **3.25**
338 **service manager**
339 responsible for managing the service instance and service topology
- 340 The service manager provides facilities to an administrator to create virtual machine instances and
341 services. The service manager also provides mechanisms for creation, monitoring, control, and reporting
342 of services.
- 343 **3.26**
344 **strong authentication**
345 authentication that requires more than the customary UserID and password. For example, biometric
346 (fingerprint), certificate (X.509), smartcard, extra questions (for example, mother's maiden name) are
347 used in addition to UserID and password to elevate the claim that the authentication is of high quality.
- 348 **3.27**
349 **virtual image**
350 virtual image is an element (often as part of a package using the [Open Virtualization Format](#)), that
351 encapsulates a workload consisting of all the code necessary to run the workload together with the
352 metadata that is necessary to configure the environment in which to run it
- 353 **3.28**
354 **workload portability**
355 provides the capability for the [cloud service consumer](#) to create a service package and then provision that
356 package in different [cloud service providers](#) without substantial modifications

357 **4 Symbols and Abbreviated Terms**

358 **4.1**

359 **API**

360 application programming interface

361 **4.2**

362 **CMDB**

363 Configuration Management Database

364 **4.3**

365 **CMIWG**

366 Cloud Management Interface Working Group

367 **4.4**

368 **CMMWG**

369 Cloud Management Model Working Group

370 **4.5**

371 **CSP**

372 cloud service provider

373 **4.6**

374 **CSPI**

375 cloud service provider interface

376 **4.7**

377 **DoS**

378 denial of service

379 **4.8**

380 **DPI**

381 deep packet inspection

382 **4.9**

383 **HTTP**

384 Hypertext Transfer Protocol

385 **4.10**

386 **IaaS**

387 Infrastructure as a Service

388 **4.11**

389 **IDS**

390 intrusion-detection system

391 **4.12**

392 **MEP**

393 Message Exchange Pattern

394 **4.13**

395 **MOF**

396 Managed Object Format

397	4.14
398	NIC
399	Network Interface Card
400	4.15
401	NIST
402	National Institute of Standards and Technology
403	4.16
404	OVF
405	Open Virtualization Format
406	4.17
407	PaaS
408	Platform as a Service
409	4.18
410	QoS
411	quality of service
412	4.19
413	RBAC
414	role-based access control
415	4.20
416	RDF
417	Resource Description Framework
418	4.21
419	REST
420	Representational State Transfer
421	4.22
422	RPC
423	Remote Procedure Call
424	4.23
425	SIP
426	Semantic Interaction Pattern
427	4.24
428	SLA
429	service-level agreement
430	4.25
431	SLO
432	service-level objective
433	4.26
434	SNIA
435	Storage Networking Industry Association

436	4.27
437	SOAP
438	Simple Object Access Protocol
439	4.28
440	SSL
441	Secure Sockets Layer
442	4.29
443	TCP
444	Transmission Control Protocol
445	4.30
446	TLS
447	Transport Layer Security
448	4.31
449	UDP
450	User Datagram Protocol
451	4.32
452	URI
453	Uniform Resource Identifier

454 **5 Introduction**

455 This document follows the definitions of cloud computing from work by the National Institute of Standards
456 and Technology ([NIST Definition of Cloud Computing](#), version 15). In part, NIST defines cloud computing
457 as "... a model for enabling convenient, on-demand network access to a shared pool of configurable
458 computing resources (for example, networks, servers, storage, applications, and services) that can be
459 rapidly provisioned and released with minimal management effort or service provider interaction."

460 NIST defines four cloud deployment models:

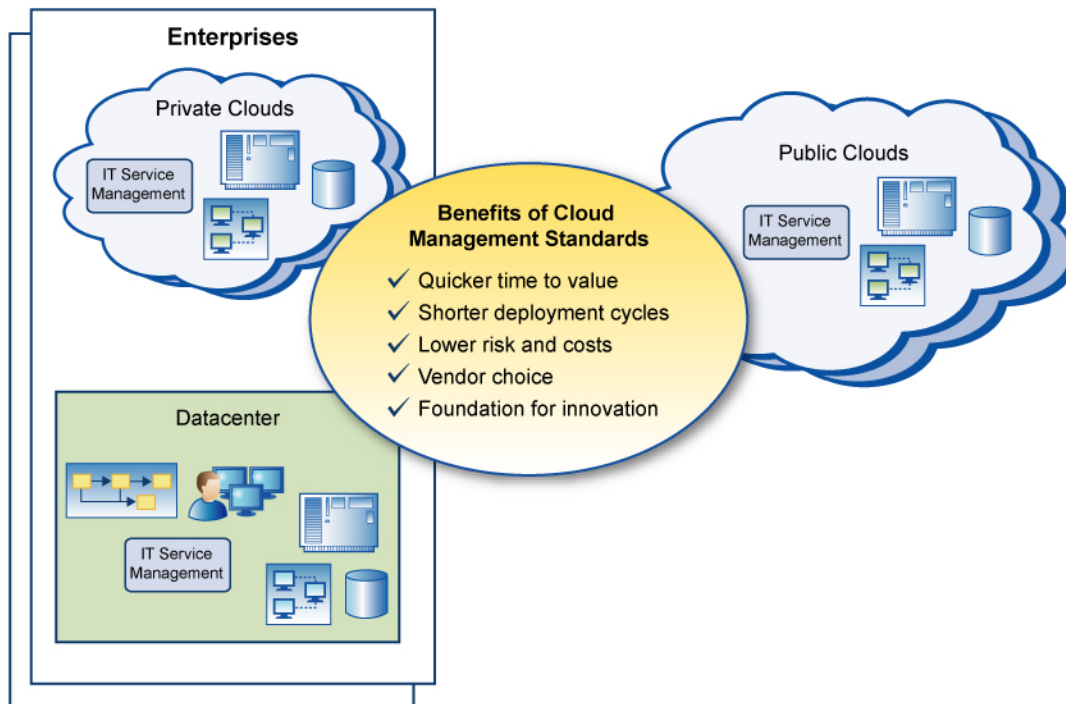
- 461 • public clouds (cloud infrastructure made available to the general public or a large industry
462 group)
- 463 • private clouds (cloud infrastructure operated solely for an organization)
- 464 • community clouds (cloud infrastructure shared by several organizations)
- 465 • hybrid clouds (cloud infrastructure that combines two or more clouds)

466 The environment under consideration by the Open Cloud Standards Incubator includes all of these
467 deployment models. The main focus of the Incubator is management aspects of Infrastructure as a
468 Service ([IaaS](#)), with some common characteristics that might be applicable to other service stacks like the
469 PaaS. These aspects include service-level agreements ([SLAs](#)), service-level objectives ([SLOs](#)), quality of
470 service (QoS), workload portability, automated provisioning, accounting, and billing.

471 The fundamental IaaS capability made available to cloud consumers is a cloud service. Examples of
472 services are computing systems, storage capacity, and networks that meet specified security and
473 performance constraints. Examples of cloud service consumers are enterprise datacenters, small
474 businesses, and other clouds.

475 Many existing and emerging standards will be relevant in cloud computing. Some of these, such as
 476 security-related standards, apply generally to distributed computing environments. Others apply directly to
 477 virtualization technologies, which are currently considered one of the important building blocks in cloud
 478 implementations. (The dynamic infrastructure enabled by technologies such as virtualization aligns well
 479 with the dynamic on-demand nature of clouds.) Examples of standards include SLA management and
 480 compliance, federated identities and authentication, and cloud interoperability and portability.

481 Figure 1 shows the scope of the Open Cloud Standards Incubator and the benefits of extending
 482 management standards.



483

484 **Figure 1 – Scope and Benefits of DMTF Open Cloud Standards Incubator**

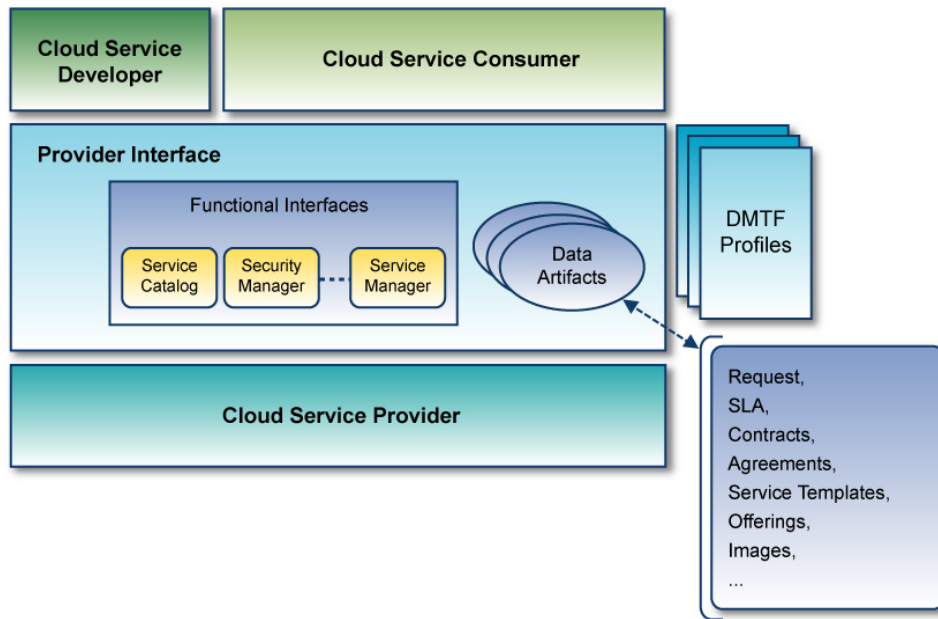
485 The Open Cloud Standards Incubator addresses the following aspects of the lifecycle of a cloud service:

- 486
- 487
- 488
- 489
- 490
- 491
- description of the cloud service in a template
 - deployment of the cloud service into a cloud
 - offering of the service to consumers
 - consumer entrance into contracts for the offering
 - provider operation and management of instances of the service
 - removal of the service offering

492 When practical, existing standards (or extensions to them) will be integrated into the recommended
 493 solution. Examples of standardization areas include resource management protocols, data artifacts,
 494 packaging formats, identity protocols, key management protocols, audit formats, compliance formats, and
 495 security mechanisms to enable interoperability.

496 **6 Reference Architecture Overview**

497 The lifecycle narrative contained in this document cites example functional interfaces that cloud
 498 consumers need to establish with the cloud service provider. This section provides a cloud service
 499 reference architecture (Figure 2) that describes key components such as actors, interfaces, data artifacts,
 500 and profiles with an indication of interrelationships among these components. This section contains a
 501 high-level discussion of the architecture, and section 9 contains more details about an expansion of the
 502 architecture. The discussion herein provides a functional description that is non-normative and profitable
 503 for study in determining architectural elements that a cloud service provider should make available.



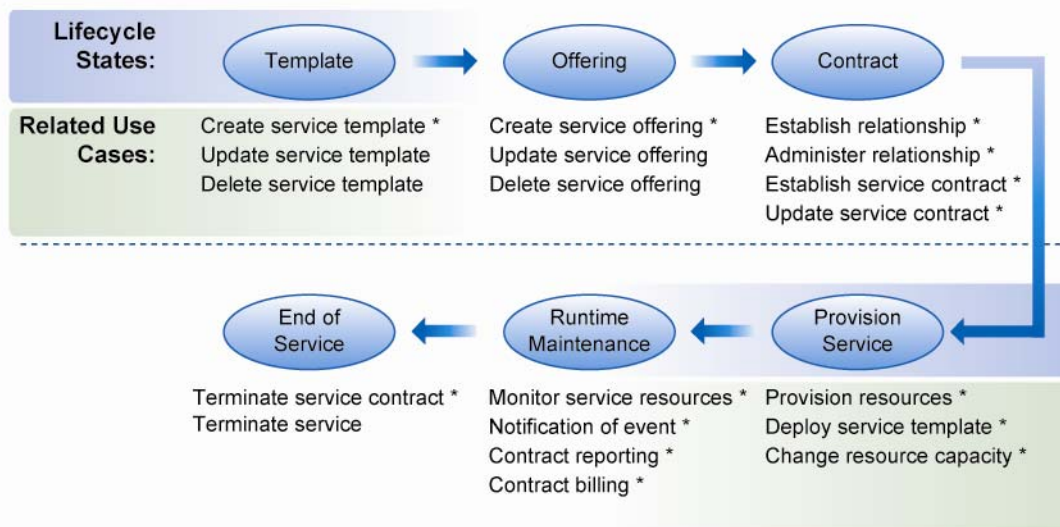
504

505 **Figure 2 – Cloud Service Reference Architecture**

506 **6.1 Service Lifecycle and Use Cases**

507 Figure 3 shows the six lifecycle states of a typical cloud service with the use cases that are most relevant
 508 to each state.

- 509 • **Template:** A developer defines the service in a template that describes the content of and
 510 interfaces to a service.
- 511 • **Offering:** A provider applies constraints, costs, and policies to a template to create an offering
 512 available for request by a consumer.
- 513 • **Contract:** A consumer and provider enter into a contract for services, including agreements on
 514 costs, SLAs, SLOs, and specific configuration options.
- 515 • **Provision Service:** A provider deploys (or modifies) a service instance per the contract with the
 516 consumer.
- 517 • **Runtime Maintenance:** A provider manages a deployed service and all its resources, including
 518 monitoring resources and notifying the consumer of key situations.
- 519 • **End of Service:** A provider halts a service instance, including reclaiming resources for
 520 redeployment to support other service.



521

522

Figure 3 – Cloud Service Lifecycle States and Use Cases

523 The use cases with asterisks in Figure 3 are described in detail in the *Use Cases and Interactions for*
 524 *Managing Clouds* white paper ([DSP-IS0103](#)).

525 6.2 Resource Models

526 Any programmatic API has an underlying resource model, whether implicit or explicit. In the IT
 527 management domain, the practice has long been to make resource models explicit and clearly separated
 528 from the protocols used to manipulate model elements. The DMTF CIM model is an example of such a
 529 protocol-independent model, for which several access protocols have been defined, supporting various
 530 interaction patterns. For large domains, such as cloud computing, there is typically not just one resource
 531 model but a series of related models. The boundaries between them are established based on
 532 considerations of reusability and complementarity. Some entities (such as *customer*, *provider*, *service*,
 533 and *policy*) are applicable independently of the type of cloud resources provided, while others (such as
 534 *host* and *disk image*) describe a specific type of cloud resources. These type-specific model entities are
 535 grouped based on the likelihood of being offered and consumed simultaneously. For example, a *host* and
 536 a *volume* are typically considered part of the same domain of cloud resources (typically called IaaS —
 537 Infrastructure as a Service). Application data storage resources (such as a SQL engine or other
 538 structured data store) would typically be grouped as part of another resource domain. In addition to
 539 deciding how to group model entities, the other key design decision is the choice of a meta-model (for
 540 example, MOF or RDF) in which the model is represented.

541 After the meta-model and models have been determined, the way the resource model is made accessible
 542 through the management interface represents the next set of design decisions. These decisions depend
 543 on how model-centric the protocol is — whether the protocol granularity corresponds exactly to the
 544 resource model granularity (resources are individually addressable and the logical protocol endpoints
 545 correspond to the model entities) or whether the model is used to describe and document the operations
 546 in the interface but the operations can take place at a different level of granularity. These
 547 addressability/granularity decisions apply to data retrieval, data setting (if direct access is permitted), and

548 invocation of operations. Other decision points include the interaction patterns and security requirements
549 and are described elsewhere in this document.

550 **6.2.1 Service Offering Model**

551 To manage cloud resources of various types, a provider/consumer resource model must exist. That
552 model describes the way in which an offering is presented and consumed. It can be abstracted from the
553 specific type of cloud resources offered. The service offering model defines entities for the service
554 consumer and the service provider. Service templates are used to describe in a general form what a
555 provider can offer, and a specific provider turns a service template into a specific offering. A catalog or
556 other mechanism is used to query and retrieve service templates and offerings. The consumer can
557 request a service instance of a service offering. Resource model elements define the lifecycle of a
558 request. The actual service delivery is modeled through domain-specific model entities, such as those in
559 the IaaS model (see 6.2.3).

560 **6.2.2 Identity Model**

561 Although cloud computing brings new use cases and threats, it does not by itself require a new identity
562 model. Existing enterprise identity models are appropriate but should be integrated, as discussed later in
563 this document. What may change is the amount of auditing kept and the precision of the user definitions
564 (individual user and organization). However, the introduction of a multi-tenant model as a cloud service
565 provider business accelerator may require more consideration. The intermingling of identity and attribute
566 information of multiple cloud service provider consumers will require careful data modeling to prevent
567 accidental disclosure.

568 **6.2.3 IaaS Model**

569 The IaaS domain is usually defined to include resources of the following types: server, whether virtual or
570 physical; storage; network connections; and IP addresses, public or private, exposed by the network
571 interfaces. In addition, the resource model for the IaaS layer may include resources that represent
572 additional features of the infrastructure, such as network features (for example, load balancing),
573 templates (for deployment), and snapshots. Data storage elements (for example, volume) have quality-of-
574 service elements of their own (for example, deduplication, encryption, backup), but these are usually
575 captured in a separate resource model (for data and storage) outside of the scope of this work.

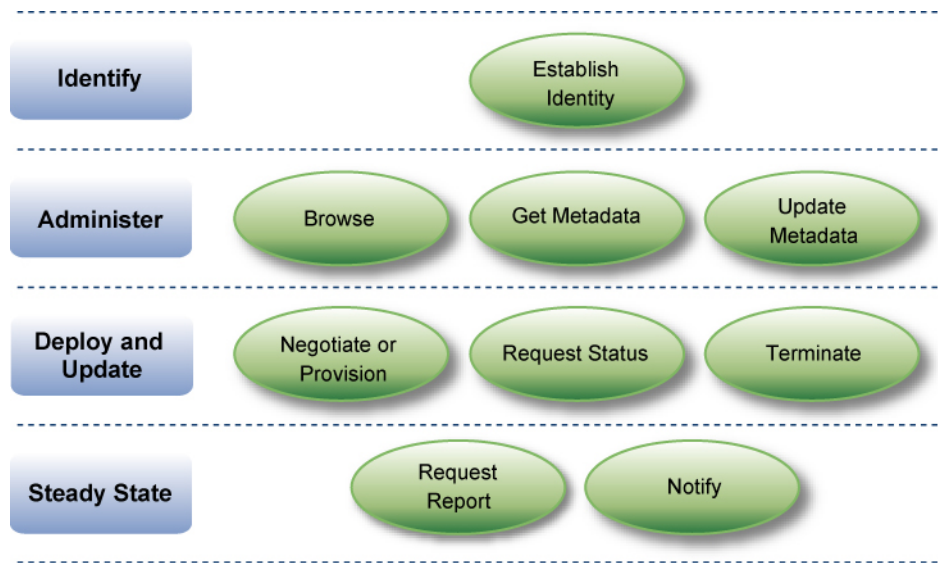
576 **6.3 Interaction Patterns**

577 An analysis of various cloud management use cases revealed that each use case could be decomposed
578 into a combination of interaction patterns. An interaction pattern consists of a sequence of messages. For
579 any interaction pattern, the specific messages may vary from use case to use case, but the messages
580 have similar characteristics at an architectural level, particularly the protocol and security considerations.

581 Figure 4 shows the interaction patterns, grouped in four broad categories.

- 582 • **Identify:** A person or entity that interacts with the cloud service provider establishes their
583 identity and receives appropriate credentials, such as a session token. An identity token may
584 also be obtained through an external identity provider that has a trust relationship with the cloud
585 service provider. Operations and data are made accessible to the connection authenticated by
586 the credentials or identity token.
- 587 • **Administer:** These patterns work with the data that describe offerings, users, and other
588 administrative metadata information needed for interactions with the cloud service provider. For
589 example, an administrator or operator may browse a list of available offerings to select one, to
590 update its metadata to configure it for a particular purpose, and to retrieve details about how to
591 access instances of a service that is part of the offering.

- 592
- 593
- 594
- 595
- 596
- 597
- 598
- 599
- 600
- **Deploy and Update:** These patterns (there are actually two types of Negotiate/Provision Resources) are used when selecting services and resources, and then making them into services. Included are any needed negotiations about the amount and type of resource, operations to provision services including the infrastructure that supports them, and tracking the status of what may be long-running operations.
 - **Steady State:** These patterns are used after services and resources have been provisioned and are in use. They include client-initiated requests such as a report request and notifications from a provider about situations that are of interest to a consumer and that may require remediation actions.



601


602

Figure 4 – Interaction Patterns

603 Figure 5 shows an example of how interaction patterns can be combined to represent a use case.

604 Provisioning a service follows the following sequence:

- 605
- 606
- 607
- 608
- 609
- 610
- 611
- 612
- 1) Establish the identity of the accessing entity.
 - 2) Browse available offerings.
 - 3) Submit and execute a request to provision a service, including supporting resources. The request may require negotiation before the consumer and provider agree to an acceptable request.
 - 4) The status of a possibly long-running provisioning process may be requested while in process.
 - 5) The consumer gets metadata about the provisioned service, such as the network addresses to access service endpoints and resources.



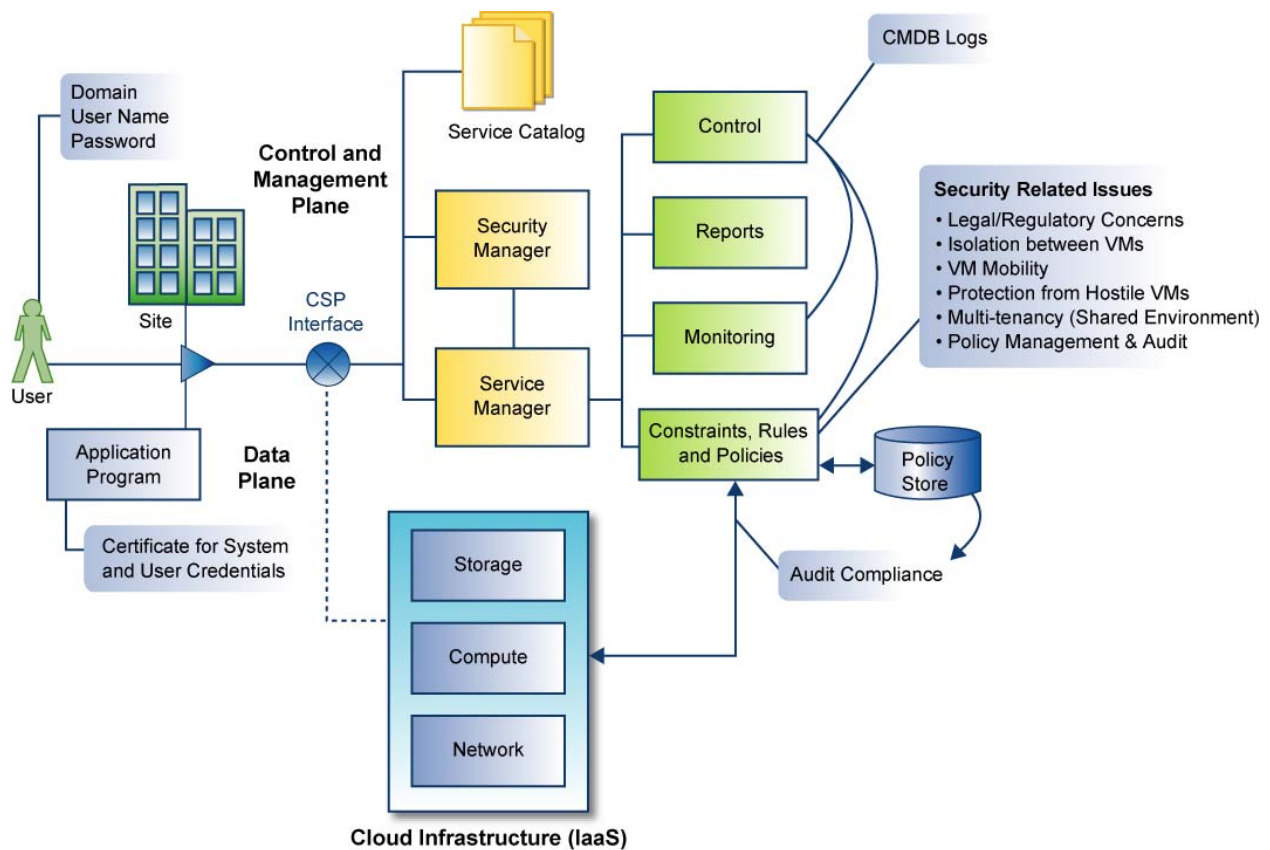
613

614

Figure 5 – Example of Interaction Patterns for a Provision Service Use Case

615 **6.4 Security Architecture**

616 In general, we are applying appropriate security primitives to the interfaces, the services, and the objects
 617 in the control and management planes. Because this work is focused on management, the security of the
 618 data plane (for example, the security of the run-time interactions between the users and the applications
 619 running on the IaaS platform) is out of the scope of the provider interface. Figure 6 shows the security
 620 context, the flow of information through the cloud service provider interface, and the objects secured.



621

622

Figure 6 – Security Context

623 The cloud service provider (CSP) interface provides access to the logical endpoints, including the [security](#)
624 [manager](#), [service manager](#), and the [service catalog](#). These endpoints provide the various services to
625 interact with service entities (such as VMs, volumes, networks, and composite applications), get audit
626 reports, and perform a host of other activities required to fulfill and maintain a cloud infrastructure.

627 The major categories of objects that are managed by the service manager are control, monitor, and report
628 objects, access to which may be controlled by role-based access control (RBAC).

629 The Constraints, Rules, and Policies objects are consumed by the cloud infrastructure, and the function of
630 the provider interface is to manage the content of these policy-related elements. The cloud infrastructure
631 (IaaS) is a “black box” to the provider interface, and how the policies are implemented is left to the cloud
632 service provider. But the fact that the various constraints, rules, and policies are being implemented is
633 verified by the audit events.

634 The two categories of actors who interact with the CSP interface are human users and application
635 programs (such as management, automatic provisioning, billing, or audit applications). The human user
636 might also be interacting through a portal interface, usually using a web browser. The portal interface will
637 be developed using the [cloud service provider interfaces](#). Both actors would be authenticated at the CSP
638 interface by the security manager or present an identity token to the security manager.

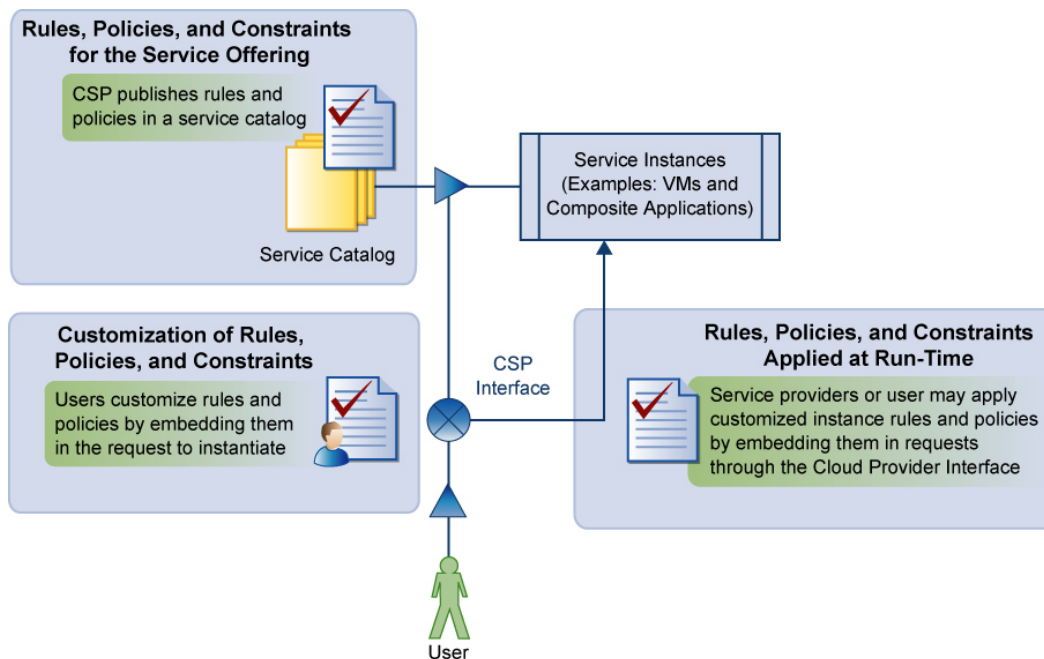
639 Traditionally, the human user uses a user name and password as credentials for authentication; however,
640 stronger mechanisms (identity federation and assertion provisioning) should be used. Commonly, an
641 application program uses certificates. However, it is deemed insecure to embed user names and
642 passwords in application programs. In this case as well, tokenized identity can be profitably used to
643 provide a higher standard of security.

644 These are common examples, but they are not prescriptive. Humans might use authenticator tokens, for
645 example, or applications might use Kerberos tickets. Appropriate mechanisms may vary in different
646 environments. Trust relationships may be employed to strengthen the authentication and authorization
647 mechanisms. Our intention is not to declare an exhaustive list of acceptable authentication mechanisms,
648 but to establish a good starting point.

649 **6.5 Role of Rules, Policies, and Constraints**

650 Policies and constraints are expected to play a major role in the management of clouds. Given the focus
651 on the interfaces for managing clouds rather than the internal details of a cloud implementation, a useful
652 abstraction is to define the desired capabilities of the cloud using policies and constraints. These are
653 interpreted to be *what* the cloud service provider should offer rather than *how* it offers it.

654 There will be several types of policies, such as service-level agreements (SLAs), service-level objectives
655 (SLOs), deployment constraints (for example, those in [OVE](#), now and proposed), data residency
656 constraints, auditability constraints, security constraints, and so on. After a consumer or developer and a
657 provider agree to a set of constraints, these constraints will govern how each (especially the provider)
658 acts. Figure 7 shows a model of the flow of constraints and rules.



659

660

Figure 7 – Constraints Flow

661 The cloud service provider publishes various relevant constraints, rules, and policies and may include
 662 them as a part of the service offering in the service catalog. These become part of the instance when a
 663 service entity (for example, a VM or a composite application) is instantiated from the template.

664 The user can customize the constraints, policies, and rules as a part of the service request to instantiate a
 665 service. Users might not be able to customize some policies and constraints from the service provider.
 666 The instantiate request will include sections to add and modify the constraints bundle. Then the instance
 667 will have an aggregate of the constraints, rules, and policies. The modified constraints have to be applied
 668 to the instance throughout its lifecycle (for example, suspend/resume or migrate).

669 After a service instance (for example, a VM or a composite application) is running, the CSP or user may
 670 request to change the policy set, within limits. It may not always be appropriate or feasible for a CSP to
 671 change the policies applied to a live service entity, which might impact the characteristics that were
 672 agreed upon before instantiation.

673 Examples of constraints are security policies. A number of different security policies can be applied to
 674 offerings in the service catalog or instances. A few simple examples follow:

- 675 • **Access control.** A service instance may be accessible only to a subset of cloud users (for
 676 example, users belonging to an enterprise or users with appropriate authorization levels). In
 677 addition, policies may specify which users are entitled to modify an instance and where and
 678 when they can deploy it.
- 679 • **Network security policies.** Such policies may specify how a service instance in a template
 680 may be connected with other service instances or resources outside the cloud. These could
 681 take the form of traffic filtering (for instance, firewall) rules and requirements for configuring
 682 traffic inspection for intrusion detection and prevention.
- 683 • **“Scope” of the security policies.** The offering may specify where instances may run; for
 684 example, they may be constrained within a particular cloud service provider’s infrastructure,
 685 enterprise zone, or geographic region.

686 **7 Requirements**

687 This section describes the high-level requirements that a cloud service provider should make available as
688 a part of the cloud service offering. The list should not be considered exhaustive but rather indicative of
689 the types of functionality that should exist in a well-rounded and exemplary cloud service provider
690 offering.

691 **7.1 Protocol Stack**

692 The following requirements are related to the protocol stack:

- 693 • The interface should support all Message Exchange Patterns (MEPs) relevant to this domain. This
694 includes one-way MEP (push and pull) and two-way MEP (synchronous and asynchronous) over a
695 request-response transport like HTTP.
- 696 • Programmatic access to various cloud provider services should be provided using industry-standard
697 mechanisms such as REST, SOAP, and so on.
- 698 • Payloads may vary dramatically in size from a few bytes to many gigabytes, and may flow between a
699 consumer or developer and a provider in either direction. Appropriate choices must be made based
700 on the situation, and the protocol should enable any of these choices.
- 701 • The stack should be able to carry additional contextual information such as authentication
702 information.
- 703 • The stack should adapt to client capability, such as availability of client address, payload size,
704 processing duration, and throughput, and use the most appropriate message pattern.
- 705 • The stack should support client- or server-initiated cancellations and restart of a large transfer.

706 **7.2 Resource Model**

707 The following requirements are related to the resource model:

- 708 • The resource model should be separated from the interfaces and mechanisms. Separating the
709 resource model from the protocol allows the consumer (and provider) applications to be isolated from
710 the protocol. As long as the applications' features are implemented based on the resource model
711 (entity types, corresponding data, associations between them, operations), they can be mapped to
712 any protocol over which the management interactions take place. This separation also allows new
713 entities, from contiguous domains, to be added to the resource model without needing to change any
714 part of the protocol design and implementation.
- 715 • The resource model should support instantiating and addressing a large number (100,000) of
716 artifacts.
- 717 • Any resources exposed at the provider interface may need to be individually addressable, at least for
718 purposes of identification, even if a resource does not support direct operations. The access to
719 various resources should have appropriate authentication (for example, discoverable resources
720 might need to expose metadata without authentication).
- 721 • Many actors may need access to resources. They will have varying entitlements, and the model
722 must support authorizations at all levels. Policies such as role-based access controls (RBACs)
723 should authorize access to and control all the resources, with visibility rules enforceable based on
724 context. Examples of the resources and capabilities whose access should be controlled include the
725 ability to create and manipulate aggregates (like virtual datacenter, virtual application stack elements
726 like VMs, network and storage, images, templates, offerings and so on), the ability to monitor
727 resources, and the ability to ask for and receive audit information about these resources.
- 728 • The metadata endpoints need not be same as the resource endpoints. For example, an audit report
729 might be a database rather than the virtual datacenter resource, or a power-on-VM interface might

730 not need to connect directly to a VM but rather to a resource representation of the VM (for example,
731 a message queue).

732 • If versioning is supported, the resource model should have the capability to create and read versions
733 of all the persistent resources, particularly those that are accessed by multiple actors and are
734 mutable.

735 • Providers must be able to expose resource metadata that may be needed by clients, such as the
736 geographical location in which data is stored.

737 **7.3 Adoptability**

738 The following requirements are related to adoptability:

739 • The service offerings should integrate with commonly found IT mechanisms, such as authentication
740 and networking.

741 • The service offerings should use easily and widely understood programming models and
742 programmer skills.

743 • The service offerings should be able to federate with other clouds, especially those that support the
744 same standards-based interfaces.

745 • Standard syntax, open APIs, and open standards should be used whenever possible.

746 **7.4 Internationalization**

747 Given that clouds are easily accessible world-wide, the interface should impose no restrictions on the
748 location and language of users, managed resources, and data. While a provider may implement
749 restrictions, the interface should not constrain the provider's choices.

750 **7.5 Rules, Constraints, and Policies**

751 The following requirements are related to rules, constraints, and policies:

752 • The cloud service provider should allow the definition of policy per consumer that will define how
753 service entities (virtual machines, network conductivity, storage, and so on) are to be treated by the
754 cloud service provider on behalf of the consumer.

755 • Whenever services are deployed by the consumer in the cloud service provider infrastructure, the
756 policies and rules imposed by the cloud service provider must be augmented by the policy of rules
757 specified by the consumer before making the images operational in the cloud service provider's
758 infrastructure. If there are any conflicts concerning policy specifications between the cloud service
759 provider and the cloud service consumer, those conflicts should be communicated to the cloud
760 service consumer for resolution.

761 • The system shall have capabilities to declaratively specify the various relevant constraints through
762 exchange of rules and policies and have the ability to verify the implementations of the said
763 constraints through audit reports. The various security threats will be mitigated by specifying policies
764 and then inspecting auditing information. How these are implemented is beyond the scope of the
765 provider interface.

766 • The rules and constraints should include:

767 – elasticity rules and constraints for compliance in terms of geographical constraints

768 – adjacency rules regarding data and processing adjacency

769 – performance affinity rules between VMs, applications, and so on

770 – anti-affinity rules for PCI compliance

- 771 • In the case of a virtualized infrastructure, there should be mechanisms to capture the hypervisor
772 security capability as well as to push policies and rules to reflect the VM security primitives. The
773 service catalog and offering should reflect these rules and constraints.
- 774 • The interface should be capable of multi-tenant security specification of items including resource and
775 data isolation, network isolation with security of virtualized networks, and so on.

776 **7.6 Cloud Management Interface Security**

777 The following requirements are related to cloud management interface security:

- 778 • Access to the cloud service provider interface endpoints and controlled resources should be secure
779 through standardized techniques. (Examples of endpoint security for cloud service providers include
780 providing protection for REST-based URLs, URL routers, and so on.)
- 781 • Programmatic and manual access to cloud provider services and content flow should be protected
782 by using industry-standard protection mechanisms such as SSL, TLS, and so on.
- 783 • The cloud service provider should assure that all access or changes to cloud services, resources,
784 and data produce auditable records regardless of success or failure, and that these audit records
785 can be compiled into a consistent audit trail that can correlate to end user- or service-initiated
786 actions. Audit records should include clear indication of any delegations of identity or authorizations.
- 787 • In cases where anonymous access is required, anonymous access should be allowed only through
788 the security manager. These cases should be recorded as exceptions and logged in the audit
789 database with clear descriptions of reasons for the exceptions and actions taken by the security
790 manager.
- 791 • The cloud service provider should provide services that establish trust relationships between itself
792 and customers or other service providers using standardized techniques; these standard
793 mechanisms include the exchange of certificates, cryptographic materials (for example, keys), and
794 root identities, as well as security policies that can be used to establish subsequent trust
795 relationships and policies.
- 796 • The cloud service provider should provide identity services that permit customers to provision and
797 manage identities that can be authenticated and authorized to access cloud services using
798 standardized mechanisms. These services should consider the need to provision identities either
799 manually or by automated means and the need to synchronize or collaborate with external identity
800 provider services.
- 801 • The cloud service provider should allow its customers to control and manage the cryptographic
802 materials and methods used to protect its applications and data, whether in motion or at rest, when
803 entering, exiting, or residing within the cloud service provider's infrastructure. The minimum
804 functionality would allow an administrator to enter this information manually, with more functionality
805 allowing secure synchronization.
- 806 • Cloud service providers should provide industry-standardized entry points into their identity provider
807 service. Such entry points allow for consumers to establish appropriate trust relationships with a
808 cloud service provider.
- 809 • The cloud service provider should provide services that can authenticate the identity of customers or
810 services and produce security information that can be consumed by cloud services to enforce
811 security policy and access control.
- 812 • Identity functionality provided by the cloud service provider may allow for the delegation of identity
813 and authorization information using standardized mechanisms such as authentication tokens. Any
814 representation of identity and authentication information, either direct or delegated, should be
815 secured and provided for. Delegation gives an account holder the ability to grant access to someone
816 who does not have an account on the service by giving them time-limited access to a specific
817 resource. Part of delegation is also auditing and keeping track of usage from these delegated

818 accesses. Further capabilities may include the ability to create an account for these accesses so that
819 the delegated client can be billed for their usage.

820 **7.7 Data and Storage**

821 The following requirements are related to data and storage:

- 822 • The cloud service provider should provide customer transparency regarding how data integrity is
823 maintained throughout the lifecycle of the data.
- 824 • Encryption and key management should use industry and government standards. The cloud service
825 provider should provide role management and separation of duties for data access controls.
826 Encryption and key management should be able to handle data isolation for multi-tenant storage and
827 separation of customer data from operational data from the service provider.
- 828 • Sensitive customer data should be encrypted in transit over the cloud service provider's internal
829 network, in addition to being encrypted at rest. The transport access to the provider interface may be
830 secured by TLS/SSL with a server certificate, and for increased security, mutual authentication at the
831 TLS layer may be optionally implemented, based on client certificates.
- 832 • In accordance with regulatory requirements, the cloud service provider should provide appropriate
833 notification to data owners when data can be (or has been) seized by a third party (for example, a
834 supply chain or the government) for content discovery.
 - 835 – Whether data can be seized by a third party could be part of the definition of the provider's
836 offering, and no notification might be meaningful there.
 - 837 – When data is seized, the decision to notify the owner of the data should be based on the
838 provider's governance.
- 839 • Data backup and recovery schemes for the cloud must be in place and effective in order to prevent
840 data loss, unwanted data overwrite, and destruction.
- 841 • Geographic data location should be exposed to the cloud consumer where appropriate. Knowing
842 where the cloud service provider will host the data and implement required measures ensures
843 compliance with local laws that restrict the cross-border flow of data.
- 844 • Data retention and secure destruction capabilities should be provided. The service provider should
845 be able to completely and effectively locate data in the cloud to be removed and provide a level of
846 assurance and evidence regarding data destruction.

847 **7.8 Logs, Event Management, Incident Response, and Notification**

848 The following requirements are related to logs, event management, incident response, and notification:

- 849 • Cloud service providers should provide incident and alert information to the customer by using
850 defined interfaces.
- 851 • Cloud service providers should maintain the confidentiality of customer incident data.
- 852 • The event management interface should have rich semantics with identifiable data sources
853 (application logs, firewall logs, IDS logs, and so on) so that the cloud consumers can make
854 inferences about their systems. The events and information should enable the cloud service
855 consumer for the verification of billing, the providing of audit events, verification of regulatory
856 compliance, the monitoring of operational characteristics, and so on.
- 857 • During the operation of service entities (for example, VMs or other services) in the cloud, the cloud
858 service provider should be monitoring the service entities to log events concerning operations, billing,
859 compliance, and so on. The cloud service provider should make this information available to the
860 cloud service consumer in accordance with the SLA in real-time, batch mode, or both.

- 861 • The event interface should also be versatile enough to address incident notification, incident
862 response, containment, and remediation.
- 863 • There should be error handling mechanisms, with error messages that are descriptive.
- 864 • Operation logs should be available in accordance with the SLA that defines the log persistence
865 duration.
- 866 • Appropriate log preservation and integrity mechanisms, such as digital signature and storage
867 devices with write-once media, should be employed.

868 **7.9 Audit, Legal, and Compliance Monitoring**

869 The following requirements are related to audit, legal, and compliance monitoring:

- 870 • Cloud service providers should include metrics and controls that customers can leverage to
871 implement their information risk management requirements.
- 872 • Cloud service providers should provide evidence on how customer security requirements are being
873 met.
- 874 • Cloud service providers should provide the interfaces to perform an external audit, where
875 appropriate.
- 876 • Cloud service providers should comply with appropriate auditing standards, such as SAS70, as
877 required by the business models and industry regulations.

878 **7.10 Security Considerations for Virtualization Technology**

879 The following requirements are related to security considerations for virtualization technology:

- 880 • Administrative access and control of virtualized operating systems should integrate with strong
881 authentication, identity management, RBAC (with separation of duty — SoD), as well as logging and
882 integrity monitoring tools.
- 883 • The cloud service provider should expose capabilities for segregating the VMs and creating security
884 zones by type of usage (for example, desktop versus server), production stage (for example,
885 development, production, and testing), location of the VM, and sensitivity of data.
- 886 • The reporting mechanism through the event management interface should provide the information
887 necessary to show data isolation and separation-of-duty in a multi-tenant environment.
- 888 • The cloud service provider should allow for the definition of virtual machines and application
889 services. The definition of such virtual machine images and services must be protected so that, if
890 required by the consumer, the images are protected from other consumers' view and reference.

891 **7.11 Portability and Interoperability for Secure Migration**

892 The following requirements are related to portability and interoperability for secure migration:

- 893 • Service entities (for example, VMs) should be able to migrate across organizational and ownership
894 boundaries (for example, between an enterprise and a service provider's IaaS infrastructure).
- 895 • In the case of a virtualized infrastructure, VM migration should address secure deprovisioning
896 (removal of the VM image after it is ported to a different location or service provider) and partial
897 migration (cloud burst: secure integration between old and new locations and service providers).
- 898 • Service providers should provide assurance on the consistency of control effectiveness,
899 management, monitoring, and reporting interfaces and their integration across old and new locations
900 and providers.

- 901 • If storage migration capabilities are provided, the service provider should have verified functionalities
- 902 for secure data transfer including encryption, access control, key management, decommissioning of
- 903 storage devices, and destruction of data after migration.

904 **8 Protocol Examples**

905 A discussion on cloud infrastructure needs to include rendering of the identified uses cases into
 906 commonly used protocols. During the analysis work undertaken to identify requirements from an
 907 infrastructure standpoint, different approaches to analysis were explored. A set of criteria roughly divided
 908 into the following different categories was used to analyze every use case:

- 909 • interaction characteristics
- 910 • resource requirements
- 911 • deployment requirements
- 912 • security requirements

913 During the analysis work for use case interactions, it was observed that a simpler set of interaction
 914 patterns were repeating in different combinations of all the use cases. It was further observed that a set of
 915 nine smaller patterns combined in various ways could render every use case that had been identified.
 916 These building blocks were termed Semantic Interaction Patterns (SIPs). As an example, Status Request
 917 is involved in many use cases and hence is considered a separate interaction pattern. Some others
 918 include Resource Negotiation and Change Notification, which get used in various use cases. The SIPs
 919 were then used in the infrastructure analysis based on the preceding categories of criteria. This section
 920 focuses on the patterns of message exchange and provides examples of how to render these patterns
 921 into commonly observed protocol styles. Using HTTP as the transport protocol, from among various styles
 922 of protocol rendering, two popular styles in IT management domain are considered: Remote Procedure
 923 Call (RPC) and Representational State Transfer (REST).

924 Some distinguishing features of each of these styles are as follows:

- 925 • RPC style assumes that each resource class provides a specific interface for manipulating the
- 926 resource. Two variants can be considered: RPC over HTTP, and SOAP RPC. Both use an XML
- 927 payload for method invocation and use the HTTP POST method to deliver it.
- 928 • REST style uses a hypertext driven paradigm with protocols such as HTTP with standard
- 929 methods GET, POST, PUT, DELETE, and others as the interfaces to a resource. Each resource
- 930 is identifiable with a URI.

931 **8.1 Message Exchange Patterns**

932 Two message exchange patterns, or MEPs, can be used to render a semantic interaction between a
 933 cloud user and cloud service: a one-way MEP, and a two-way MEP. Furthermore, to reflect various
 934 connectivity constraints or user profiles, it is useful to define “execution modes” for these MEPs that can
 935 be associated with specific requirements or environment constraints. This section illustrates a few
 936 execution modes as applicable to one-way and two-way MEPs. The execution modes pertain to the two
 937 actions “push” and “pull” that the sender and the receiver may take in order to complete the interaction. It
 938 should be noted that *sender* and *receiver* are generic terms and that the entity sending back the response
 939 may not be the same addressable entity that received the request.

940 The protocol options show a few variants in using HTTP. A protocol solution has several aspects: RPC
 941 versus Document styles, RPC versus REST, plain HTTP with REST-style versus SOAP-based solutions.
 942 The focus here is on REST versus RPC. Some SOAP-based solutions that do not rely on RPC (for
 943 example, document/literal SOAP style) are not described here.

944 **8.1.1 One-Way MEP**

945 Execution modes for one-way MEP, as illustrated in Figure 8, are as follows:

- 946 • The **One-Way/Push MEP**, or “push” mode, where the sender takes the initiative of sending the
947 business message to the receiver (for example, over an HTTP request). This mode assumes
948 that the receiver has an IP address. An example would be a change notification interaction
949 where the notified party — the user — is addressable.
 - 950 – Typical steps by a client and service for a one-way/push MEP using RPC:
 - 951 1) The client (here the sender) issues an HTTP POST request with an RPC payload that
952 contains the object instance ID, method call, and parameters.
 - 953 2) The service sends an HTTP response — a status code that indicates success or
954 failure. No RPC payload is expected.
 - 955 – Typical steps by a client and service for a one-way/push MEP using REST:
 - 956 1) The client (here the sender) issues an HTTP command (PUT, POST, or DELETE) to a
957 URI that represents the particular resource.
 - 958 2) The service responds with an HTTP success or failure response.

959 See Annex A for protocol examples using HTTP.

- 960 • The **One-Way/Pull MEP**, or “pull” mode, where the receiver takes the initiative of pulling the
961 business message from the sender (for example, by initiating an HTTP request that contains a
962 pull signal for a particular queue or mailbox). The response is sent over the HTTP response.
963 This mode allows receivers that do not have an address to receive business messages.

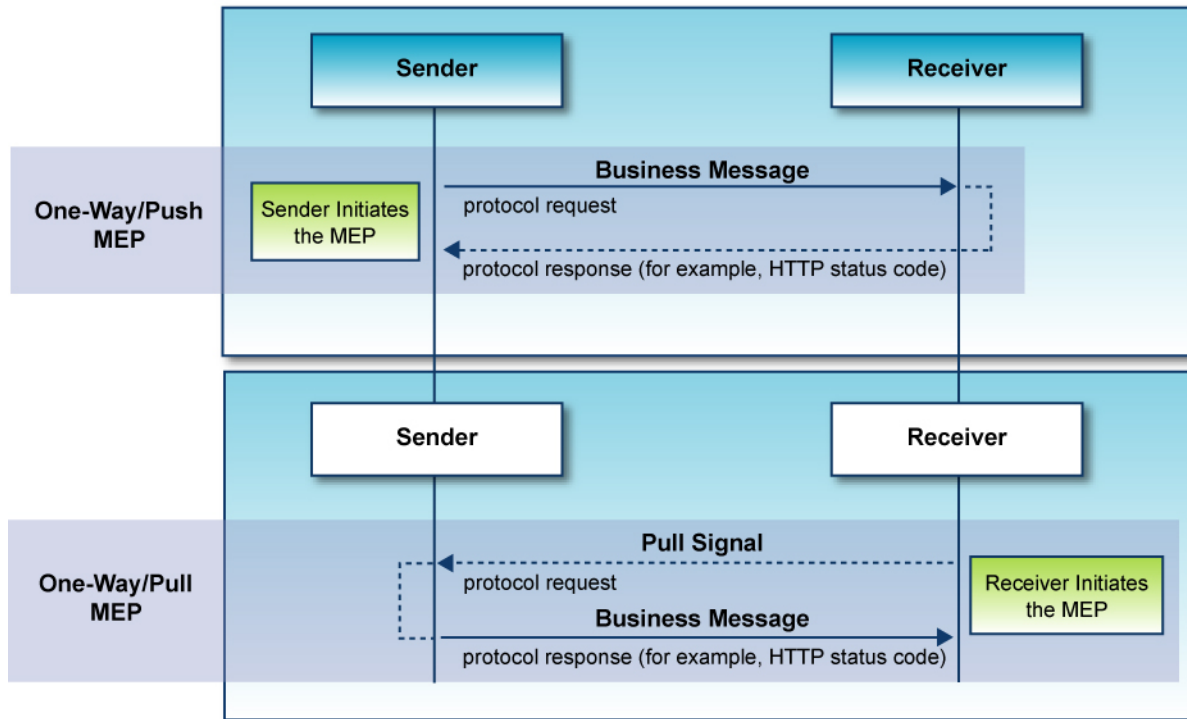
964 An example would be a change notification interaction in which the notified party — the user —
965 is not addressable and must pull notifications that are obtained on the protocol back-channel
966 (for example, HTTP response).

967 In the pull mode, the receiver might periodically send several pull signals before a successful
968 MEP is complete (for instance, including a business message on the response leg). These pull
969 signals are typically scheduled automatically within the underlying messaging layer (for
970 example, using WS-MakeConnection [[Web Services Make Connection \(WS-MakeConnection\)](#)
971 [Version 1.1](#)], issuing an HTTP GET, or using a PullRequest signal in [OASIS ebXML Messaging](#)
972 [Services Version 3.0](#). Such pull signals contain only the absolute minimal information necessary
973 to pull (for example, a queue ID, a channel ID) and no other business information.

974 Here it is assumed that the client can provide a client side URI prior to or as a part of the call
975 that is issued. In other pull variants, a queue ID or mailbox ID is provided to the client for use in
976 all its subsequent pulls from the service.

- 977 – Typical steps by a client and service for a one-way/pull MEP using RPC:
 - 978 1) Optionally, the client receives a notification with an RPC payload on the URL provided
979 with an event ID.
 - 980 2) The client issues an HTTP GET or POST with an event ID (or a queue ID) as a
981 parameter to get the details.
 - 982 3) The responding party sends back a payload about resource information that
983 corresponds to the event ID.
- 984 – Typical steps by a client and service for a one-way/pull MEP using REST:
 - 985 1) Optionally, the client receives a notification on the URI of the event.
 - 986 2) The client issues an HTTP GET of the event's URI.
 - 987 3) The responding party sends back a payload about the resource identified by the URI.

988 See Annex A for protocol examples using HTTP.



989

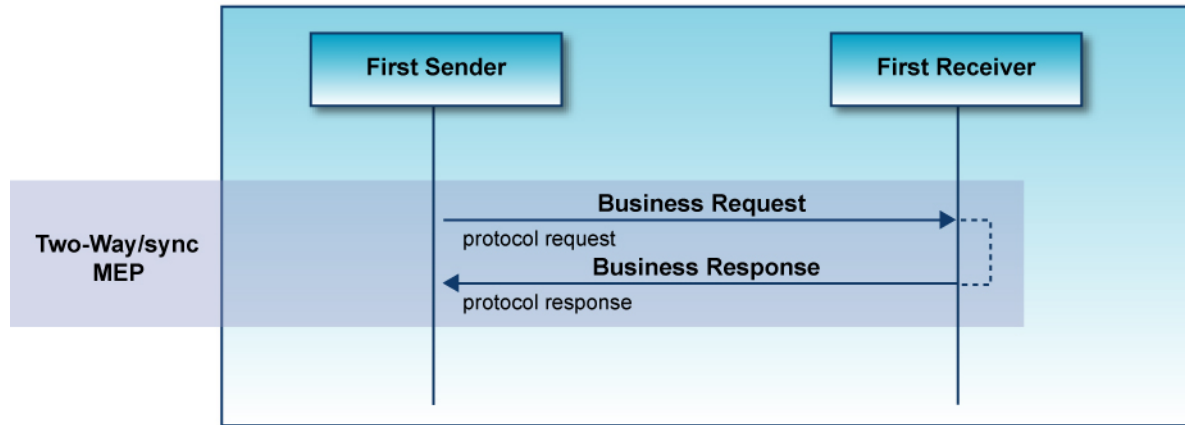
990

Figure 8 – One-Way MEP Execution Modes

991 **8.1.2 Two-Way MEP**

992 Execution modes for two-way MEP are as follows:

- 993 • In the **Two-Way/Sync MEP**, or “sync” mode, the sender takes the initiative of sending the
 994 business request to the receiver and gets the business response over the same connection (in
 995 case of a request-response transport such as HTTP). This mode does not require the sender to
 996 be addressable. An example would be requesting status interaction, assuming that the
 997 generation of the status information is quick and can accommodate a single HTTP request-
 998 response exchange. Figure 9 shows an illustration.



999

1000

Figure 9 – Two-Way/Sync MEP Execution Mode

1001

- Typical steps by a client and service for a two-way/sync MEP using RPC:

1002

(This is quite similar to the one-way/push MEP example. The primary difference is that this execution mode contains an RPC response payload.)

1003

1004

- 1) The client issues an HTTP POST request with an RPC payload that contains request data such as the object instance ID, the method call, and parameters.

1005

1006

- 2) The client receives an HTTP response with the RPC payload if the method succeeds, or an HTTP error code failure if the method fails.

1007

1008

- Typical steps by a client and service for a two-way/sync MEP using REST:

1009

- 1) The client issues an HTTP command (GET, PUT, POST, or DELETE) to a URI that represents the particular resource. The client expects to receive a business response (which is different from a one-way/push MEP).

1010

1011

1012

- 2) The client receives an HTTP response with a payload that represents resource data if the method succeeds.

1013

1014

See Annex A for protocol examples using HTTP.

1015

- The **Two-Way/Push-and-Push MEP** is an asynchronous mode, where the business request and the business response are sent over different connections (for example, both pushed over a separate HTTP connection). Both sender and receiver are supposed to be addressable. This mode is useful when the response requires some computation time. An example would be a simple resource-negotiation interaction, in which the cloud service may take some time to allocate requested resources, beyond the delay tolerated for a single HTTP request-response. In that case, the confirmation is sent over a callback. Figure 10 shows an illustration.

1016

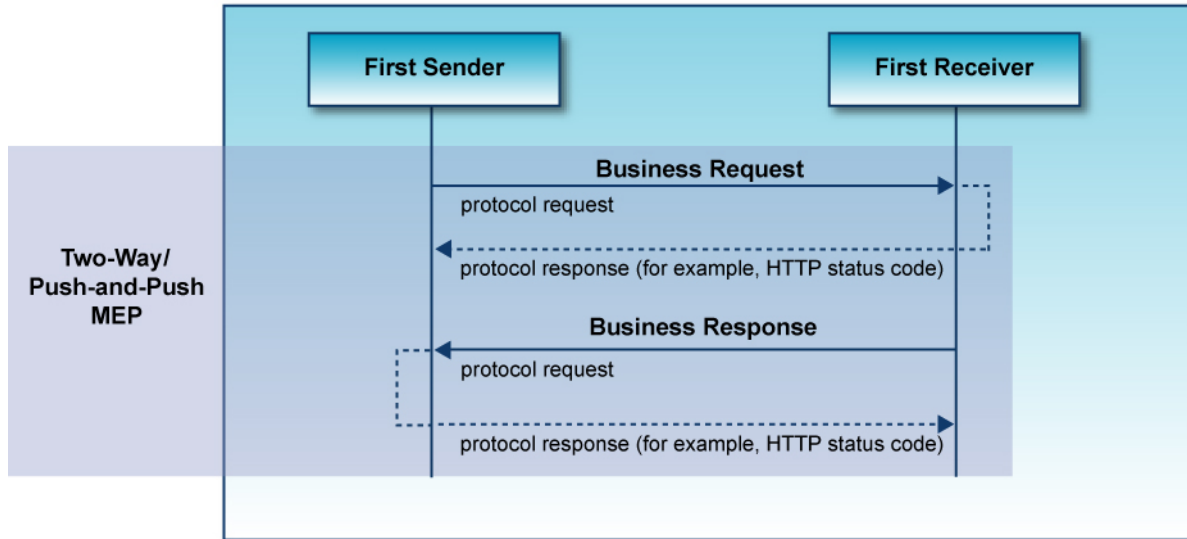
1017

1018

1019

1020

1021



1022

1023

Figure 10 – Two-Way/Push-and-Push MEP Execution Mode

1024

- Typical steps by a client and service for a two-way/push-and-push MEP using RPC:

1025

1) The client issues an HTTP POST request with an RPC payload that contains request data including the client's addressable endpoint.

1026

1027

2) Optionally, the service responds with an identifier that represents the work that is proceeding in response to the request.

1028

1029

3) Sometime later, the service issues a call-back (for example, through HTTP POST) on the client URL with the payload that contains the response to the initial Method Call.

1030

1031

- Typical steps by a client and service for a two-way/push-and-push MEP using REST:

1032

1) The client issues an HTTP command to the URI that identifies the resource.

1033

2) Optionally, the service responds with a URI from which the client can obtain status information.

1034

1035

3) Sometime later, the service issues a call-back (generally through HTTP POST) on the URI that the client registered.

1036

1037

See Annex A for protocol examples using HTTP.

1038

- The **Two-Way/Push-and-Pull MEP** is an asynchronous mode, where the sender takes the initiative of both exchanges: pushing the business request and pulling the business response. The advantages over the "push-and-push" mode are that the sender is not required to have an addressable endpoint and can decide when it is ready to transfer the response. An example would be a simple resource-negotiation interaction, in which the cloud user is not addressable, and needs to pull the confirmation message later. Figure 11 shows an illustration.

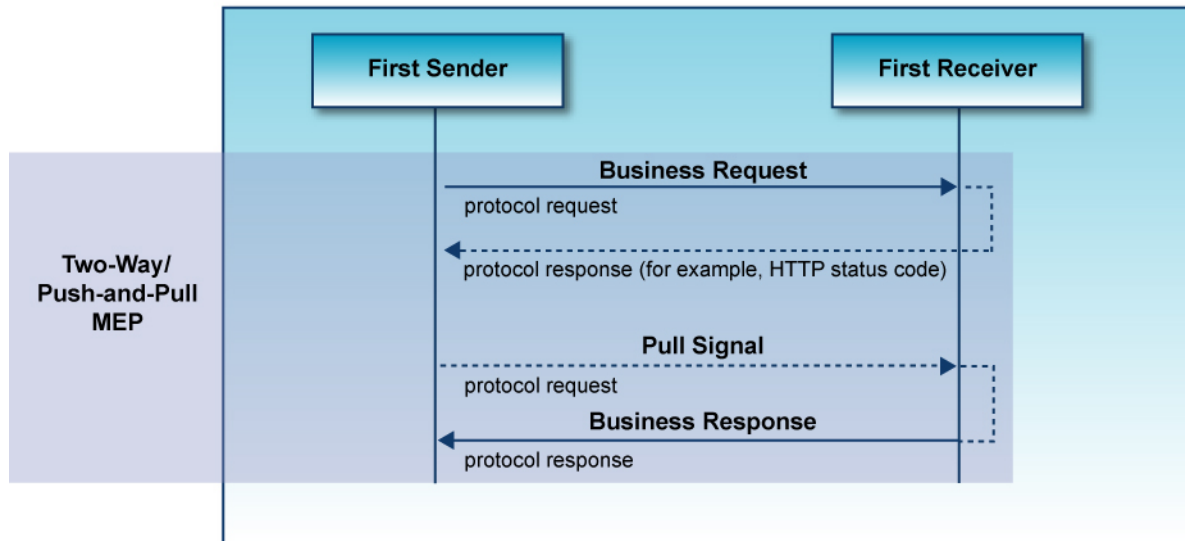
1039

1040

1041

1042

1043



1044

1045

Figure 11 – Two-Way/Push-and-Pull MEP Execution Mode

1046

- Typical steps by a client and service for a two-way/push-and-pull MEP using RPC:

1047

- 1) The client (here “First Sender”) issues an HTTP POST request with a payload that contains the object instance ID, the method call, and parameters.

1048

1049

- 2) Optionally, the service (here “First Receiver”) responds with a payload that contains an identifier to be used by the client in a subsequent query, unless the ID is already agreed upon.

1050

1051

1052

- 3) The client issues an HTTP POST request with a payload that contains the ID received in the response to the initial request.

1053

1054

- 4) The service sends back the business RPC payload over the HTTP response. It sends back an empty payload if no message is available for pulling.

1055

1056

- Typical steps by a client and service for a two-way/push-and-pull MEP using REST:

1057

- 1) The client (here “First Sender”) issues an HTTP request to the URI that identifies the resource.

1058

1059

- 2) Optionally, the service (here “First Receiver”) responds with a payload that contains a URI to be used by the client in a subsequent request.

1060

1061

- 3) The client issues an HTTP GET request to the URI.

1062

- 4) The service sends back the business payload (for example, a notification) over the HTTP response, or the service may simply return an HTTP respond code without any payload.

1063

1064

1065

See Annex A for protocol examples using HTTP.

1066 8.2 Infrastructural Aspects Affecting the Choice of MEP and Its Execution Mode

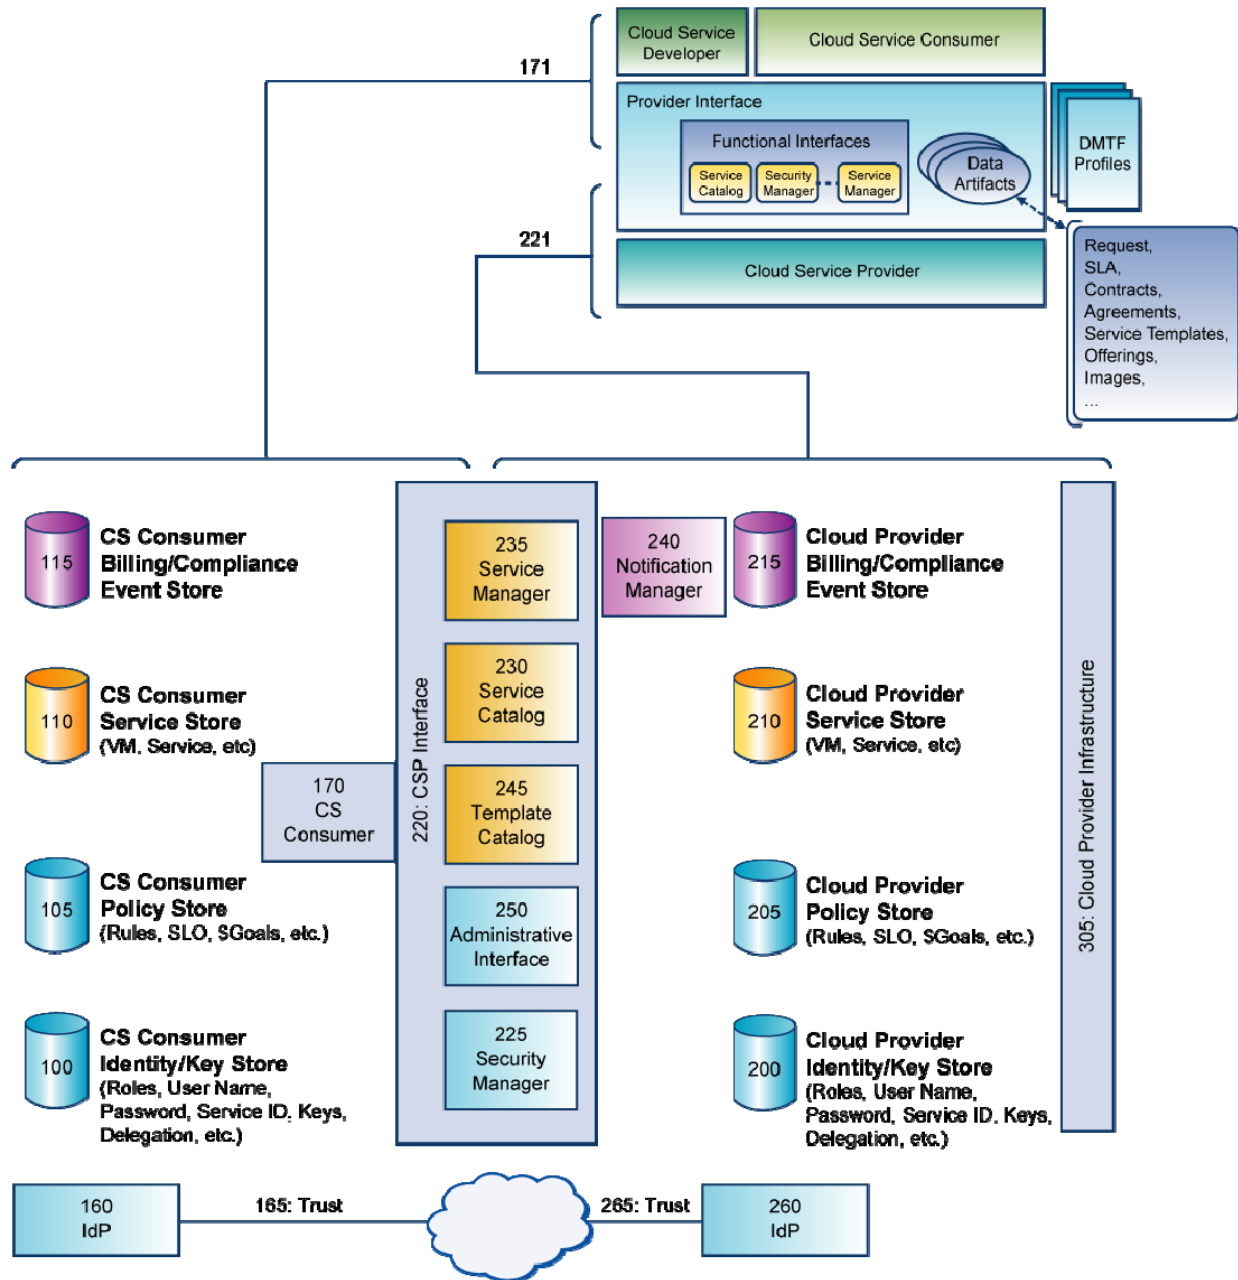
1067 Some infrastructural aspects have more of a bearing on the choice of the MEP than others.

- 1068 • The payload size informs whether an interaction is likely to be short or long, and
1069 correspondingly whether a synchronous or an asynchronous pattern of interaction is preferable.
- 1070 • The amount of time that is likely to be taken by the backend in the processing affects the choice
1071 of execution mode. The time that it takes to execute an operation may be completely
1072 independent of the payload that it takes to execute it. A long delay may prohibit use of a two-
1073 way/sync MEP because it may exceed the timing tolerated for an HTTP request-response.
1074 However, if the response time is constrained to be short (for example, by SLA), then a two-
1075 way/sync MEP can be used.
- 1076 • Client addressability affects the choice of execution mode. Only if the client is addressable by
1077 the cloud can the two-way/push-and-push MEP (for instance, with “callback”) or the one-
1078 way/push MEP directed to the client be used. If the client is *not* addressable, it will have to pull
1079 all asynchronous responses to its requests.
- 1080 • Finally, the ability to browse or enumerate a large collection also has a bearing on which MEPs
1081 to use for browse operations. This is a variation of the first bullet, though there may be different
1082 processing time characteristics because of namespace enumeration operations.

1083 9 Cloud Ecosystem

1084 This section expands on the reference architecture described earlier in this document and rationalizes the
1085 architecture with data artifacts described in the *Use Cases and Interactions for Managing Clouds* white
1086 paper ([DSP-IS0103](#)). Familiarity with [DSP-IS0103](#) is required to fully understand this section’s contents.
1087 By understanding this section, the reader can expect to come to an understanding of the architectural
1088 interactions and functional capabilities that would be needed by an offering from a cloud service provider.
1089 The information contained here should be considered referential only and non-normative.

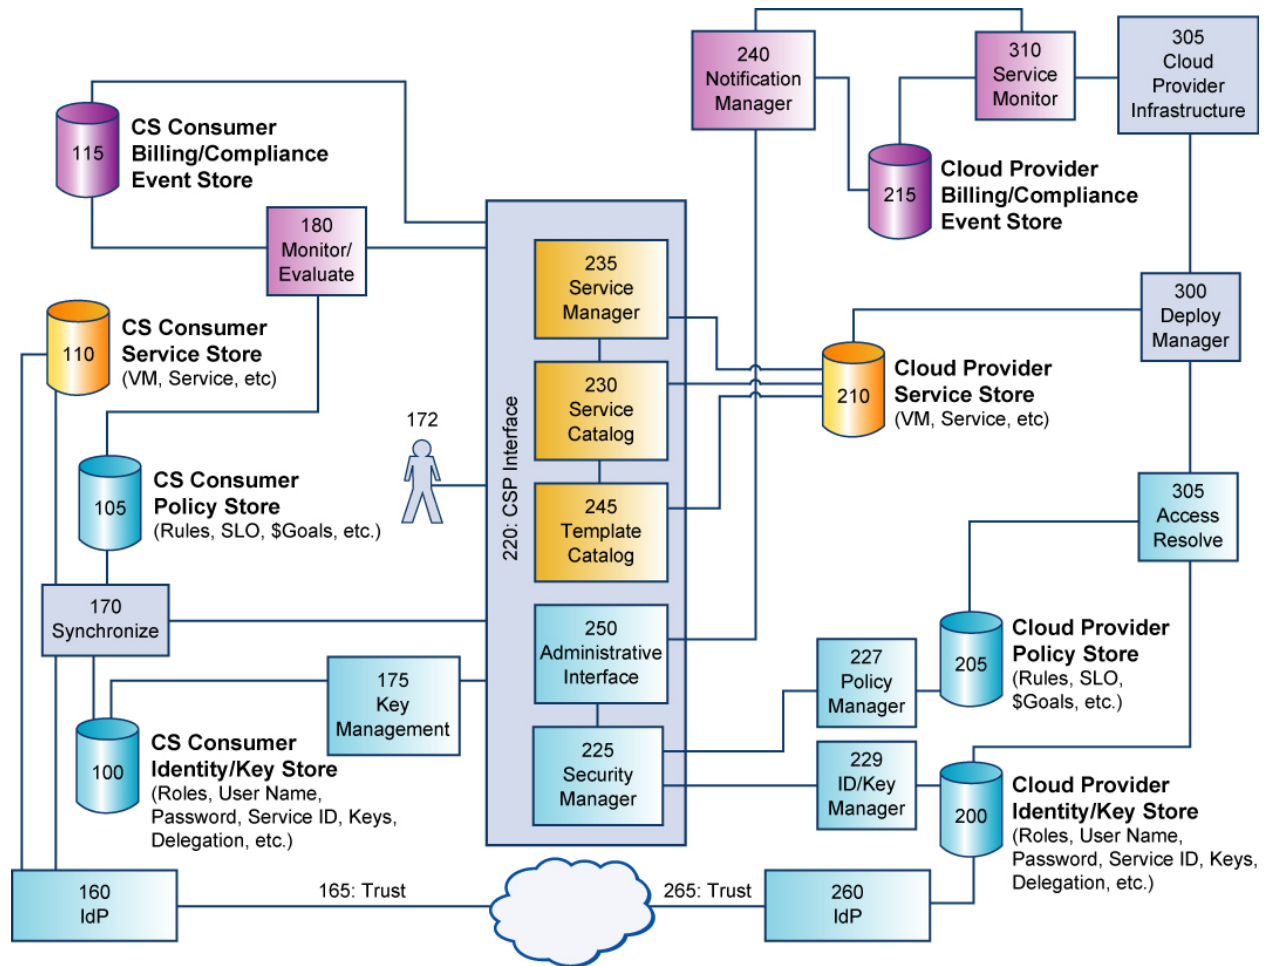
1090 Figure 12 is provided to help the reader understand the relationship of the architecture presented in
1091 *Interoperable Clouds* ([DSP-IS0101](#)) and the more detailed architectural diagram and discussion
1092 presented in Figure 13. In the area labeled 221, Figure 12 shows that the provider interface offered by the
1093 cloud service provider is represented by elements labeled 225–250, and the cloud service provider is
1094 represented by the areas labeled 200–215, 240, and 305. The area labeled 171 shows that the cloud
1095 service consumer is represented by the areas labeled 100 through 170. The identity provider relationships
1096 are shown in the areas labeled 160, 165, 265, and 260. The function of the identity provider service (IdP)
1097 at 160 and 260 together with the trust relationship at 165 and 265 is shown to help the reader make the
1098 transition to the more detailed diagram in Figure 13.



1099

1100

Figure 12 – High-Level Architecture



1101

1102

Figure 13 – Expanded Architecture

1103 The discussion of Figure 13 begins by referencing the data artifacts that are managed by the interfaces
 1104 225, 230, 235, 245, and 250, along with their supporting data stores that are represented by 200, 205,
 1105 210, 215, 100, 105, 110, and 115. The interfaces are responsible for managing the lifecycle of the data
 1106 artifacts described in the *Use Cases and Interactions for Managing Clouds* white paper ([DSP-IS0103](#)).

1107 **9.1 Architectural Impact on Data Artifacts**

1108 The Service Manager [235] manages the following data artifacts described in [DSP-IS0103](#):

- 1109
- 1110 • service instance
 - 1111 • service topology
 - 1112 • service topology item
 - 1113 • service topology relationship

1114 The Service Catalog [230] manages the following data artifacts described in [DSP-IS0103](#):

- 1115
- 1116 • service offering
 - 1117 • service contract

- 1116 • service request
- 1117 The Template Catalog [245] manages the following data artifact described in [DSP-IS0103](#):
- 1118 • service template
- 1119 The Security Manager [225] manages the following data artifacts described in [DSP-IS0103](#):
- 1120 • UserIdentityInformation
- 1121 • CredentialsToken
- 1122 The Administrative Interface [250] manages the following data artifacts described in [DSP-IS0103](#):
- 1123 • service reporting information
- 1124 • service event log
- 1125 • service notification
- 1126 • All other data artifacts not mentioned above

1127 **9.2 Cloud Service Provider Identity/Key Store [200]**

1128 The cloud service provider identity/key store [200] provides the architectural abstraction for the cloud
 1129 service provider to store and provide local identities, roles, cryptographic keys, delegation specifications,
 1130 and so on. The data artifacts in the cloud service provider identity/key store [200] can be provided by the
 1131 cloud service provider but are more commonly provided by a cloud consumer administrator [172], which
 1132 manages the user names, passwords, roles, and so on for personnel associated with the consumer's
 1133 business to use cloud resources. Using an administrator to provision the cloud service provider
 1134 identity/key store [200] is labor intensive and has been found to be less than desirable by cloud
 1135 consumers in today's market.

1136 Another source of data artifacts in the cloud service provider identity/key store [200] is a synchronize
 1137 function at 170 sponsored by the cloud consumer. This function synchronizes user names, passwords,
 1138 roles, delegation specifications, and so on from an identity/key store on the consumer's premises [100].
 1139 This mechanism provides for much more rapid provisioning and deprovisioning of authorized cloud users
 1140 and role and delegation information. However, unless the synchronize function at 170 can provide a
 1141 password other than the password used by personnel to authenticate to the consumer's business and a
 1142 single sign-on mechanism, the user name and password that allows access to the consumer's internal
 1143 data systems is exposed. The alternative is to provide a different user name and password at the cloud
 1144 service provider identity/key store [200], but this has its own issues with personnel remembering the user
 1145 name and password and those credentials being exposed through "sticky notes" on the monitor.

1146 For consumer businesses that have access to an identity provider [160], a trust relationship [165, 265]
 1147 can be established with the cloud service provider's identity provider [260], and provisioning and
 1148 deprovisioning of user names, passwords, and so on can be performed from attributes contained within
 1149 the identity assertion produced by the identity provider [160]. This method still suffers from the possibility
 1150 of exposing sensitive credentials at the cloud service provider identity/key store [200] or requiring multiple
 1151 credentials for the user (although, the identity provider [160] can broker the use of multiple user names
 1152 and passwords and therefore lift the burden from the user).

1153 The most secure mechanism for providing consumer identity to the cloud service provider is through the
 1154 identity provider [260], but rather than populating the cloud service provider identity/key store [200] with
 1155 user names and passwords, the identity token from the identity provider [160] is used directly during
 1156 access resolution at the cloud service provider infrastructure [305] without ever storing any user
 1157 credential information in the cloud service provider identity/key store [200]. An alternative mechanism is to
 1158 store the identity assertions in the cloud service provider identity/key store [200] and use those in lieu of
 1159 user names and passwords as long as the identity assertion remains valid (for example, does not expire).

1160 In the case of using the synchronize function [170] and cloud consumer administrator [172], the cloud
 1161 service provider will need to provide the functions of an administrative interface [250], which acts as the
 1162 interface or API for marshaling information through the security manager [225] to the ID/key manager
 1163 [229], which is responsible for keeping the cloud service provider identity/key store [200] up to date.
 1164 These functions (together with all other functions shown in Figure 13) are architectural only and represent
 1165 functions to be performed rather than implementation mechanisms.

1166 Another function affecting the cloud service provider identity/key store [200] is the maintenance of
 1167 cryptographic keys on behalf of the consumer. These keys can be provided by either an administrator
 1168 [172] or a key management mechanism [175]. In either case, the administrative interface [250] provides
 1169 the functions necessary to marshal the key information to the cloud service provider identity/key store
 1170 [200] through the security manager [225]. Storing cryptographic materials at the cloud service provider
 1171 identity/key store [200] risks the possibility of cryptographic material exposure to the outside world. As
 1172 with identity, key material can also be passed through identity tokens, which can provide cryptographic
 1173 material without storing it locally at the cloud service provider’s premises (for example, 160, 165, 260,
 1174 265).

1175 To illustrate security policies in the cloud interface, Figure B-1 in Annex B describes security policies from
 1176 the point of view of the organization. Table 1 provides a cross reference of items from Figure B-1 in
 1177 Annex B to Figure 13:

1178 **Table 1 – Cross-Mapping of Elements in Figure B-1 with Elements in Figure 13**

Annex B Policy Model (Figure B-1)	References to Figure 13
Items 120,121, 122, and 123	Items 110, 210, and 245
Item 124	Items 230 and 235
Item 126	Item 105
Item 127	Item 205
Item 125	Items 105 and 205
Item 255	Item 305
Item 310	Item 310
Item 315	Item 215
Items 320 and 325	Item 180

1179 **9.3 Cloud Service Provider Policy Store [205]**

1180 The cloud service provider policy store [205] provides a repository for policies that govern the use of the
 1181 cloud service provider’s infrastructure [305], service-level objectives (SLOs), monetary goals involving
 1182 negotiated charges with the consumer (\$Goal), firewall settings for consumer deployments, and so on.
 1183 The cloud service provider policy store [205] is maintained by the policy manager [227], which,
 1184 functionally, is associated with the security manager [225]. Commonly, the cloud service provider
 1185 specifies its policies and that allows consumer administrators [172] to select the policies that they will use
 1186 and provide some limited modification of those policies. Manual maintenance of consumer policies at the
 1187 cloud service provider through the cloud service provider policy store [205] is labor intensive and is not
 1188 scalable. This situation argues for a consumer policy store [105] wherein the consumer manages the
 1189 policies, rules, and so on that are synchronized with the cloud service provider’s policy store [205] (these
 1190 policies are generally maintained by the cloud consumer as a part of their CMDB). Synchronizing policy
 1191 from the consumer sites to the cloud service provider site is especially advantageous when the consumer
 1192 is hosting a private cloud and wants to be able to use the same policies for private and public cloud
 1193 deployments. Such a synchronization of policies requires arbitration between policy specifications at the

1194 cloud service provider's policy store [205] and policy specifications at the consumer policy store [105]
1195 provided by 170, 250, 225, and 227.

1196 **9.4 Cloud Service Provider Service Store [210]**

1197 The cloud service provider service store [210] provides the cloud storage for virtual machines and
1198 services. Services are defined in this paper as a collection of virtual machines that have been configured
1199 to provide a specific service. For example, an e-mail system may comprise an SMTP Gateway, IMAP4
1200 server, POP3 server, storage, DNS server, and so on. The service may contain rules for load balancing
1201 and load shedding so that the service reacts appropriately to traffic decreases, increases, and
1202 imbalances.

1203 Traditionally, the administrator [172] uses the service catalog [230] or template catalog [245] to discover
1204 what services and virtual machines are already configured in the cloud service provider service store
1205 [210]. Configuration of a virtual machine or service is performed through the service manager [235], which
1206 also provides facilities for the administrator [172] to define new virtual machines and services.

1207 Today's use of private clouds makes it advantageous to allow the consumer's business to define its own
1208 services and virtual machines that could be used in both a private cloud and multiple public clouds. This
1209 functionality is depicted in the architecture as the CS Consumer Service Store [110]. The synchronization
1210 function shown at 170 provides the mechanisms for synchronizing the consumer service store [110] with
1211 provider service store [210] through functions and services provided by the service manager [235].
1212 Synchronization can both add and remove virtual machines and services, and as this happens the service
1213 catalog [230] is updated. The service catalog [230] and the service manager [235] are both protected by
1214 identity so that virtual machines and services are accessible only by the consumer audiences for which
1215 they are intended (whether public or restricted to a membership list). The architecture also depicts a
1216 relationship between the consumer service store [110] and the identity provider [160]. As identity and
1217 policy-applied-by-identity become increasingly important to regulatory compliance and other audit
1218 concerns, it will be important for virtual machines and services to be protected by both digital signatures
1219 and digital identities. These digital identities, and possibly the cryptographic materials, will be provided by
1220 the identity provider [160]. Because of the trust relationship [165, 265], virtual machines and services thus
1221 tagged with identity are recognizable by the cloud service provider's infrastructure (specifically by the
1222 deploy manager [300]) so that policy and access can be regulated by the identity and tampering can be
1223 recognized.

1224 **9.5 Cloud Service Provider Billing/Compliance Event Store [215]**

1225 A critical facility in the architecture supporting cloud service provider functionality is the cloud service
1226 provider Billing/Compliance Event Store [215]. As virtual machines and services operate within the cloud
1227 service provider infrastructure [305], it is critical that monitoring take place at the service monitor [310] to
1228 populate the cloud service provider billing/compliance event store [215] with the appropriate billing events
1229 and regulatory compliance events. These events are used by the cloud service provider to provide billing
1230 materials at the end of the month, provide audit materials for the customer to verify the charges from the
1231 cloud service provider, and provide an audit trail concerning regulatory compliance.

1232 Consuming businesses are asking for ways to track their use of cloud resources and to watch for audit
1233 events from cloud processes and resources in real-time. This is shown in the architecture at CS
1234 Consumer Billing/Compliance Event Store [115]. This event store is provided either through a pull or push
1235 model with the billing and compliance events monitored by the cloud service provider through the
1236 notification manager [240], and is managed through the administrative interface [250]. The notification
1237 manager [240] can also provide real-time event flow to the monitor/evaluate function [180], which can
1238 allow the consumer's business to monitor cloud evaluation and capture possible regulatory compliance
1239 deviations in real-time. The monitor/evaluate function [180] can also provide for forensic evaluations from
1240 the data contained within the event store [115]. In all of the evaluations by the monitor/evaluate function
1241 [180], the determination of whether events may be causing compliance deviation or exceeding cost goals
1242 is determined by the policy stored in the consumer policy store [105]. Likewise, a service-level objective

1243 (SLO) can be used to verify that the service-level agreement (SLA) to which the consumer and the cloud
1244 service provider agreed is being honored.

1245 **9.6 Cloud Service Provider Deployment [300]**

1246 Finally, the architecture depicts deploying virtual machines and services into the cloud service provider
1247 infrastructure [305]. This is performed by first resolving access issues [305] and then matching the access
1248 request with the services by the deploy manager [300]. It is the responsibility of the deploy manager [300]
1249 to ensure that services and virtual machine utilization are managed according to policy and identity. The
1250 policy includes both the internal policy by the cloud service provider and policies provided by the
1251 consumer. Services are deployed into the cloud service provider infrastructure [305] and finally
1252 deprovisioned from the cloud service provider infrastructure according to either a manual request by the
1253 administrator [172] or some procedural or policy specification through the administrative interface [250].

1254 **10 Conclusion and Next Steps**

1255 The current phase of the Open Cloud Standards Incubator is concluded. The work of the incubator has
1256 resulted in three documents outlining the use cases, data artifacts, and architectural considerations for
1257 IaaS clouds. This output will be handed over to the appropriate DMTF working groups and alliance
1258 partners so that they can develop standards while the Open Cloud Standards Incubator considers other
1259 potential areas to tackle.

1260 **10.1 Standards Development Steps**

1261 To take the output of the incubator and produce DMTF standard specifications, one or more working
1262 groups are needed to model the domain and design the details of the infrastructure and the interfaces.
1263 Deliberations in the incubator resulted in the request to create two working groups:

- 1264 • the Cloud Management Model Working Group (CMMWG) under the Platform Management
1265 Subcommittee)
- 1266 • the Cloud Management Interface Working Group (CMIWG) under the Infrastructure
1267 Subcommittee)

1268 After the DMTF Board approves the appropriate working groups, their charters will be posted at
1269 <http://www.dmtf.org/about/committees/>.

1270 **10.2 Important Extensions to Be Considered**

1271 This section describes further work that needs to be done to improve the flexibility, interoperability, and
1272 ease of acquisition of a DMTF cloud. The work that is remaining is itemized in this section, and it is
1273 expected that the two working groups will address these work items in some way as part of their
1274 deliberations.

1275 **10.2.1 Interface Extensibility**

1276 The incubator's interface is "narrower" than available cloud interfaces. For example, what will the
1277 consumer do if it needs both? Use them side by side? Have a DMTF interface that is extendable so more
1278 detailed interfaces can be exposed by an implementation that is under the umbrella of the DMTF
1279 interface?

1280 **10.2.2 Template Extensibility**

1281 It may be necessary to provide the means of adding new types of entities to catalogs, and composing
1282 existing entities in a catalog into another entity. For example, if a consumer wants to add a firewall in front
1283 of an existing service, how is it represented? The interface needs to have features such as copying an

1284 existing template and modifying parts of it, or downloading a template into a metadata file (for example,
1285 OVF) on the client end, modifying it, and uploading it back to the cloud.

1286 **10.2.3 Management Interface Granularity**

1287 The cloud interfaces provided by the implementations in the industry broadly fall into three groups:

- 1288 • an administrative group responsible for aspects like the overall relationship, service contract, and
1289 catalog management
- 1290 • a workload management group that deals with workloads, involving service instances and their
1291 lifecycle
- 1292 • a resource management group that is responsible for management of individual artifacts such as
1293 a VM or logical server, NICs, and storage volumes

1294 The incubator interface represents the administrative group and, to some extent, the workload
1295 management group. It does not have any artifacts to represent the resource management group.

1296 **10.2.4 Security Model Granularity**

1297 The security model may need further clarification such as mapping the main security categories (identity
1298 and access management, data protection, security event management, and software/platform/
1299 infrastructure [compliance] assurance) to the cloud reference architecture. Furthermore, it may be
1300 necessary to define specific interfaces for cloud users to have the option of more granular visibility of and
1301 control over how their security policies and requirements are met by the service providers.

1302 **10.3 Integration with Alliance Partner Frameworks**

1303 The Open Cloud Standards Incubator looks to alliance partners to provide domain-specific
1304 implementations and expects the DMTF cloud implementations to interoperate with them. An example is
1305 a storage cloud standard being worked on by SNIA. Activities need to be undertaken to provide guidance
1306 on connecting the DMTF cloud interfaces and those created by alliance partners such as SNIA. Work
1307 may be needed both from a resource-model-compatibility perspective as well as from a protocol
1308 perspective.

1309 **11 Future of the Cloud Incubator**

1310 With regard to the Open Cloud Standards Incubator, two areas of further work have been identified: cloud
1311 federation and Platform as a Service (PaaS). It is recommended that the incubator charter be extended to
1312 do further work in these areas.

Annex A (informative)

1313
1314
1315
1316
1317

Message Exchange Protocol Examples

1318 This annex provides examples of how various Message Exchange Patterns (MEPs) introduced in 9.1 can
1319 be rendered using XML-RPC and REST-oriented mechanisms.

1320 A.1 One-Way/Push MEP: RPC Example

1321 A.1.1 Client Request: Set User Age

```
1322 POST /MyWebService.Contoso.com HTTP/1.1
1323 Content-Type: text/xml
1324 Content-length:
1325 <?xml version="1.0"?>
1326 <methodCall>
1327 <methodName>Users.SetUserAge</methodName>
1328 <params>
1329 <param>
1330 <name><string>Brian Young</string></name>
1331 </param>
1332 <param>
1333 <value><i4>23</i4></value>
1334 </param>
1335 </params>
1336 </methodCall>
```

1337 A.1.2 Service Response

```
1338 HTTP/1.1 200 OK
1339 Host: contoso.com
1340 Content-Type: text/xml
1341 Content-length:0
```

1342 A.2 One-Way/Push MEP: REST Example

1343 A.2.1 Client Request: Set User Age

```
1344 PUT /uri_of_brain_young
1345 Host: contoso.com
1346 Content-Type : application/vnd.cloud.com.User+xml
1347 Content-Length: xxxx
1348 <xml>
1349 <user>
1350 <age value="23"/>
1351 </user>
1352 </xml>
```

1353 A.2.2 Service Response

```
1354 HTTP/1.1 200 OK
1355 Host: contoso.com
1356 Content-Type: text/xml
1357 Content-length:0
```

1358 A.3 One-Way/Pull MEP: RPC Example

1359 The one-way/pull MEP is different from a request-response (two-way) synchronous MEP in the sense that
 1360 the “pulled” message is not the result of a query with a specific selector that may vary for each MEP
 1361 execution; it is more akin to reading a queue of messages. Typically, a pull message only knows about a
 1362 channel or queue ID (the equivalent of a mailbox ID). A variant may rely on a notified “event ID” that the
 1363 client must use later only once.

1364 In order to pull, the client must have prior knowledge of this event ID or channel ID. The event ID or
 1365 channel ID could be provided dynamically by the service calling a particular client method. In these
 1366 examples, it is assumed that the client can provide a client-side URL prior to or as a part of the call that is
 1367 issued.

1368 A.3.1 Service Initiated Notification

```
1369 <!-- In this variant, an “event ID” is initially notified to the client by the
1370 service. The event ID is to be used for subsequent pulling. ---!>
1371
1372 POST /MyAddress.Client.com HTTP/1.1
1373 Content-Type: text/xml
1374 Content-length:
1375
1376 <?xml version="1.0">
1377   <methodCall>
1378     <methodName>notifyEvent</methodName>
1379     <params>
1380       <param><eventID><i4>1234</i4></eventID></param>
1381       <param>
1382         <eventDescription><string>"disk low"</string></eventDescription>
1383       </param>
1384     </params>
1385   </methodCall>
```

1386 A.3.2 Client Response

```
1387 HTTP/1.1 200 OK
1388 Host: client.com
1389 Content-Type: text/xml
1390 Content-length:0
```

1391 A.3.3 Client Request: Pull Event

```
1392 <!-- Pull the Event Request using a one-way/pull method. ---!>
1393 POST /MyWebService.Contoso.com HTTP/1.1
1394 Content-Type: text/xml
1395 Content-length:
```

```

1396
1397 <?xml version="1.0"?>
1398 <methodCall><methodName>eventGetInfo</methodName></methodCall>
1399   <params>
1400     <param>
1401       <eventID><i4>123456</i4></eventID>
1402     </param>
1403   </params>
1404 </methodCall>
1405 <!-- Pull the Event Response ---!>

```

1406 **A.3.4 Service Response**

```

1407 HTTP/1.1 200 OK
1408 Host: contoso.com
1409 Content-Type: text/xml
1410 Content-length:
1411
1412 <?xml version="1.0"?>
1413 <methodResponse>
1414   <params>
1415     <param>
1416       <maxDisk><kilobytes>512000000</kilobytes></maxDisk>
1417     </param>
1418     <param>
1419       <usedDisk><kilobytes>501000000</kilobytes></usedDisk>
1420     </param>
1421   </params>
1422 </methodResponse>

```

1423 **A.4 One-Way Pull MEP: REST Example**

1424 **A.4.1 Service Initiated Notification**

```

1425 POST /MyAddress.Client.com/Events
1426 Content-Type : application/vnd.cloud.com.Event+xml
1427 <event>
1428   <eventURI>/Events/1234</eventURI>
1429   <eventid>1234</eventid>
1430   <eventMessage>disk low</eventMessage>
1431 </event>

```

1432 **A.4.2 Client Response**

```

1433 HTTP/1.1 200 OK
1434 Host:Client.com
1435 Content-length:0

```

1436 **A.4.3 Client Request**

```

1437 GET /MyWebService.Contoso.com/Events/1234
1438 Accept :application/vnd.cloud.com/Event+xml

```

1439 A.4.4 Service Response

```
1440 HTTP/1.1 200 OK
1441 Host: contoso.com
1442 Content-Type : application/vnd.cloud.com/Event+xml
1443 Content-Length:xxx
1444
1445 <?xml version="1.0"?>
1446     <event>
1447         <eventid>1234</eventid>
1448         <eventMessage>disk low</eventMessage>
1449         <eventDetail>
1450             For Volumn abc: 501000000 out of 512000000 are used
1451         </eventDetail>
1452     </event>
```

1453 A.5 Two-Way/Sync MEP: RPC Example

1454 A.5.1 Client Request: Enumerate Users with Filter

```
1455 POST /MyWebService.Contoso.com HTTP/1.1
1456 Content-Type: text/xml
1457 Content-length:
1458 <?xml version="1.0"?>
1459     <methodCall>
1460         <methodName>Users.EnumerateUsers</methodName>
1461         <params>
1462             <param>
1463                 <filter><string>*B*</string>
1464             </param>
1465         </params>
1466     </methodCall>
```

1467 A.5.2 Service Response

```
1468 HTTP/1.1 200 OK
1469 Host: contoso.com
1470 Content-Type: text/xml
1471 Content-length:
1472
1473 <?xml version="1.0"?>
1474     <methodResponse>
1475         <params>
1476             <param>
1477                 <name><string>Betty White</string></name>
1478             </param>
1479             <param>
1480                 <name><string>Brian Young</string></name>
1481             </param>
1482         </params>
1483     </methodCall>
```

1484 A.6 Two-Way/Sync MEP: REST Example

1485 A.6.1 Client Request: Enumerate Users with Filter

```
1486 GET /Users?name=[*B*]
1487 Host:MyWebService.Contoso.com
1488 Accept:application/vnd.cloud.com/Users+xml
```

1489 A.6.2 Service Response

```
1490 HTTP/1.1 200 OK
1491 Host: contoso.com
1492 Content-Type : applicationapplication/vnd.cloud.com/Users+xml
1493 Content-Length:
1494
1495 <?xml version="1.0"?>
1496   <users>
1497     <user>
1498       <name>Betty White</name>
1499     </user>
1500     <user>
1501       <name>Brian Young</name>
1502     </user>
1503   </users>
```

1504 A.7 Two-Way/Push-and-Push MEP: RPC Example

```
1505 <!----- Issue API call to format volume ----!>
```

1506 A.7.1 Client Request: Format Volume

```
1507 POST /MyWebService.Contoso.com HTTP/1.1
1508 <?xml version="1.0">
1509   <methodCall>
1510     <methodName>formatVolume</methodName>
1511     <params>
1512       <param>
1513         <volumeID><string>"some volume"</string></volumeID>
1514       </param>
1515       <param>
1516         <callbackURL><string>"http://MyAddress.Client.com"</callbackURL>
1517       </param>
1518     </params>
1519   </methodCall>
```

1520 A.7.2 Service Initial Response

```
1521 HTTP/1.1 200 OK
1522 Host: contoso.com
1523 Content-Type:text/xml
1524 Content-Length:
1525
```

```

1526 <?xml version="1.0"?>
1527   <params>
1528     <param>
1529       <jobID><i4>1234</i4></jobID>
1530     </param>
1531     <param>
1532       <eventDescription>
1533         <string>"volume format Initiated"</string>
1534       </eventDescription>
1535     </param>
1536   </params>

```

1537 **A.7.3 Service Initiated Progress Notification**

```

1538 POST /MyAddress.Client.com
1539 Content-Type: text/xml
1540 Content-Length:
1541
1542 <?xml version="1.0"?>
1543   <methodResponse>
1544     <params>
1545       <param><jobID><i4>1234</i4></jobID></param>
1546       <param><maxDisk><kilobytes>512000000</kilobytes></maxDisk></param>
1547       <param><usedDisk><kilobytes>0</kilobytes></usedDisk></param>
1548     </params>
1549   </methodResponse>

```

1550 **A.7.4 Client Response to the Service**

```

1551 HTTP/1.1 200
1552 Host: contoso.com
1553 Content-Length:0

```

1554 **A.8 Two-Way/Push-and-Push MEP: REST Example**

1555 **A.8.1 Client Request: Format Volume**

1556 The client has already registered the callback URI before sending the request.

```

1557 PUT /Volumes/MyVolumeuri
1558 Host: MyWebService.Contoso.com
1559 Content-Type: application/vnd.cloud.com.Volume+xml
1560 Content-Length: xxxx
1561
1562 <xml>
1563   <volume>
1564     <format>desired_format</format>
1565   </volume>
1566 </xml>

```

1567 **A.8.2 Service Initial Response**

```
1568 HTTP/1.1 200 OK
1569 Host:Contoso.com
1570 Content-type: text/xml
1571 Content-length:0
```

1572 **A.8.3 Service Initiated Progress Notification**

```
1573 POST /MyAddress.client.com/
1574 Content-type : applicationapplication/vnd.cloud.com.Event+xml
1575 Content-length :
1576 < ?xml version="1.0" ?>
1577 <xml>
1578     <event>
1579         <resource
1580             uri="/Volumes/MyVolumeuri"
1581             type="application/vnd.cloud.com.Volume+xml"
1582             status="READY"/>
1583         <detail>Reformatting the volume from aaa to bbb</detail>
1584     </event>
1585 </xml>
```

1586 **A.8.4 Client Response**

```
1587 HTTP/1.1 200 OK
1588 Host: Client.com
1589 Content-length:0
```

1590 The client can then use the resource URI returned in the event to GET the details of the volume.

1591 **A.9 Two-Way/Push-and-Pull MEP: RPC Example**

1592 **A.9.1 Client Request: Format Volume**

```
1593 <!-- Issue API call to format volume ----!>
1594
1595 POST /MyWebService.Contoso.com HTTP/1.1
1596 Content-Type: text/xml
1597 Content-length:
1598
1599 <?xml version="1.0">
1600     <methodCall>
1601         <methodName>formatVolume</methodName>
1602         <params>
1603             <param><volumeID><string>"MyVolumeID"</string></volumeID></param>
1604         </params>
1605     </methodCall>
```


1606 A.9.2 Service Response

```
1607 HTTP/1.1 200 OK
1608 Host:Contoso.com
1609 Content-Type: text/xml
1610 Content-Length:
1611 <!-- receive jobID to poll --!>
1612
1613 <?xml version="1.0"?>
1614   <methodResponse>
1615     <params>
1616       <param><jobID><i4>1234</i4></jobID></param>
1617     </params>
1618   </methodResponse>
```

1619 A.9.3 Client Request: Check Job Status

```
1620 POST /MyWebService.Contoso.com HTTP/1.1
1621 Content-Type: text/xml
1622 Content-length:
1623 <!--PULL with jobID as the parameter ---!>
1624
1625 <?xml version="1.0"?>
1626   <methodCall>
1627     <methodName>pollJob</methodName>
1628     <params>
1629       <param><jobID><i4>1234</i4></jobID></param>
1630     </params>
1631   </methodCall>
```

1632 A.9.4 Service Response

```
1633 <!-- get back progress information ---!>
1634 HTTP/1.1 200 OK
1635 Host: contoso.com
1636 Content-Type: text/xml
1637 Content-Length:
1638
1639 <!xml version="1.0"?>
1640 <methodResponse>
1641   <params>
1642     <param><jobID><i4>1234</i4></jobID></param>
1643     <param><percentComplete><i4>74</i4></param>
1644   </params>
1645 </methodResponse>
```

1646 **A.10 Two-Way/Push-and-Pull MEP: REST Example**

1647 **A.10.1 Client Request: Format Volume**

```

1648 PUT /Volumes/MyVolumeuri
1649 Host: MyWebService.Contoso.com
1650 Content-Type : application/vnd.cloud.com.Volume+xml
1651 Accept : application/vnd.cloud.com.Volumn+xml, application/vnd.cloud.com.Event+xml
1652 Content-Length: xxxx
1653
1654 <xml>
1655   <volume>
1656     <format>desired_format</format>
1657   </volume>
1658 </xml>

```

1659 **A.10.2 Service Response**

```

1660 HTTP/1.1 200 OK
1661 Content-type : application/vnd.cloud.com.Event+xml
1662 Location: /Events/ThisEventUri
1663 Content-Length:xxxx
1664
1665 <?xml version="1.0"?>
1666   <xml>
1667     <event uri="/Events/ThisEventUri">
1668       <resource
1669         uri="/Volumes/MyVolumeuri"
1670         type="application/vnd.cloud.com.Volume+xml"
1671         status="REFORMATTING" percentage="0"/>
1672       <detail>reformatting volume from aaa to bbb</detail>
1673     </event>
1674   </xml>

```

1675 **A.10.3 Client Request**

```

1676 <!------- Poll for the status of the job ---!>
1677
1678 GET /Events/ThisEventUri
1679 Host: MyWebService.Contoso.com
1680 Accept : application/vnd.cloud.com.Event+xml,

```

1681 **A.10.4 Service Response**

```

1682 http/1.1 200 OK
1683 Content- type: application/vnd.cloud.com.Event+xml
1684 Location: /Events/ThisEventUri
1685 Content-Length:xxxx
1686 <?xml version="1.0"?>
1687   <xml>
1688     <event uri="/Events/ThisEventUri">

```

```
1689     <resource
1690         uri="/Volumes/MyVolumeuri"
1691         type="application/vnd.cloud.com.Volume+xml"
1692         status="REFORMATTING" percentage="74"/>
1693     <detail>reformatting volume from aaa to bbb</detail>
1694 </event>
1695 </xml>
```

Annex B
(informative)

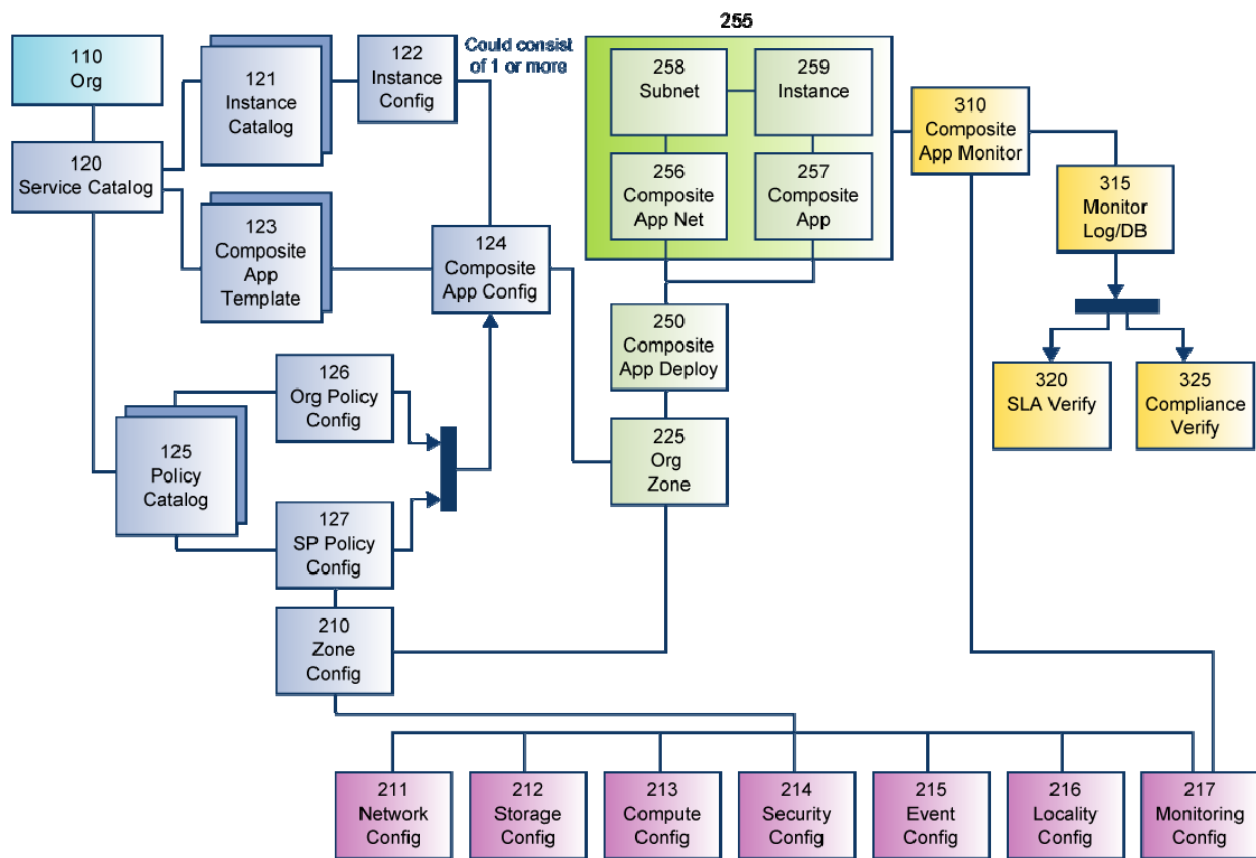
Policy Discussion

1696
1697
1698
1699
1700

1701 This annex describes what a policy model for clouds might look like. It is not intended to be prescriptive,
1702 but rather exemplary.

B.1 Policy Model

1704 The policy model could be represented by Figure B-1.



1705
1706

Figure B-1 – Policy Model

B.1.1 Policy Model from the Organization’s Perspective

1708 Consider an organization [110] that wants to use the cloud at a service provider. The organization [110] is
1709 constrained by its contract with the service provider to only the services in the service provider’s catalog
1710 that the organization approves (a part of the contract between the organization and the service provider).
1711 The service provider’s catalog consists of an instance catalog [121], a composite application template

1712 [123], and a policy catalog [125]. The instance catalog [121] consists of a collection of the individual
1713 instances [122], and creates a VM repository.

1714 The composite application template [123] consists of the collection of instances and templates defined by
1715 the organization for use. These templates could be uploaded by an API, created out of individual instance
1716 groupings, deployed in the service provider's cloud, or uploaded manually.

1717 The policy catalog [125] consists of the organization's policy configurations [126] and the service
1718 provider's policy configurations [127]. The service provider's policy configurations [127] consist of the
1719 service provider policies that govern the use and deployment of the infrastructure of the IaaS (for
1720 example, network [211], storage [212], compute [213], security [214], event [215], locality [216], and
1721 monitoring [217] configurations).

1722 The organization now constructs a repository of services at the service catalog [120], which comprises
1723 instance specifications (for example, pointers to instances [122]) and configuration and operational
1724 stipulations (for example, SLO specifications). The organization previously specified policy [126] that
1725 defines the organization's governance expectations concerning services when they become operational.
1726 The collection of 120, 121, 123, and 126 can be considered part of the organization's CMDB.

1727 **B.1.2 Policy Model from the Service Provider's Perspective**

1728 The service provider specifies the operational characteristics and configurations of each of the zones that
1729 compose the overall IaaS cloud. This is represented by zone configuration [210], which consists of a
1730 collection of zone configurations comprising network configuration [211], storage configuration [212],
1731 compute configuration [213], security perimeter configuration [214], event configuration [215], locality
1732 configuration [216], and monitoring configuration [217]. Each of these provides detailed hardware and
1733 interconnect configurations necessary to specify the infrastructure policy enforcement configurations.

1734 The service provider's policy configuration [127] will provide the governance for the collection of zone
1735 configurations [210] (because there are multiple zones, there will be multiple policies governing the zones
1736 and zone interactions). The collection at 210 and 127 can be considered part of the service provider's
1737 CMDB.

1738 The composite application configuration process [124] uses information at 123, 126, 127, and 210 to
1739 provide a coherent definition of a composite application's zone configuration [225]. The composite
1740 application configuration process [124] must provide a normalization of the policies of 126 and 127 so that
1741 the expectations of both the organization [110] and the service provider are enforced. If these policy
1742 expectations cannot be realized, the specific composite application configuration process [124] will not be
1743 created, and an exception will be asserted to both the organization [110] and the service provider.

1744 The resulting composite application's zone configuration [225] can be deployed [250] into one or more
1745 service provider's infrastructures [255]. These infrastructures include 256, 257, 258, and 259. Note that
1746 "media" is implied but not specified in Figure B-1.

1747 The resulting deployment is monitored [310], providing a log [315]. Processes to verify the SLA through
1748 SLOs [320] and operational compliance to regulatory agencies [325] are provided. SLA verification [320]
1749 and compliance verification [325] may be operational in the organization's data center, in the cloud
1750 service provider, or a combination of both.

1751 An example of the production of a coherent definition of a deployable composite application configuration
1752 could be as follows:

- 1753 1) The service provider has noted in the policy configuration [127] that all 10 GB network
1754 interconnects are available only to platinum customers.
- 1755 2) The organization under consideration is a customer that has purchased a standard performance
1756 package from the service provider. The standard performance package specifies that only
1757 100 MB and 1 GB network interconnects are available to the organization.

- 1758 3) When the composite application configuration is generated, all 10 GB network interconnects are
1759 removed from any of the zone configurations that are available to the organization. In this way,
1760 although the organization is participating in a zone that has 10 GB network conductivity, that
1761 conductivity is not available to the organization because any 10 GB network interface cards are
1762 not available to instances being deployed by the organization.
- 1763 4) If the organization upgrades its customer status to that of a platinum customer, the 10 GB
1764 network interconnects are not culled from the zone configuration, which allows any composite
1765 applications to use the higher bandwidth without changes to the application templates,
1766 instances, and so on.
- 1767 5) Likewise, if the contract period for being a platinum customer expires and the customer reverts
1768 to standard status, the enhanced network bandwidth is not available.
- 1769 6) The organization may have specified in its organization policy [126] that certain services must
1770 be connected to 10 GB networks. If the customer type is that of "standard," a conflict exists
1771 between the policy of the service provider and the policy of the organization. This conflict is
1772 noted to both the service provider and the organization for resolution.
- 1773 7) Because policy between the organization and the service provider is being normalized during
1774 composite application configuration [124] and verified during composite application deployment
1775 [250], capabilities can be added and subtracted from any given deployment because of policy
1776 resolution based on configuration (for example, customer type).

1777 **B.1.3 Policy Examples**

1778 Table B-1 – lists examples of policy types grouped by area.

1779 **Table B-1 – Policy Examples**

Policy Area	Examples
User and identity	Identity: Federation Authentication Authorization
Operational attributes	Access control User profiles Key/Certificate management Patching Perimeter Virtualization security Monitoring, incident response, and notification Secure workload migration
Legal and compliance	Electronic discovery Data jurisdiction Data retention and destruction Data location Privacy
Network and connectivity	Perimeter firewall template (to enable organizations to specify the network perimeter firewall rules): <ul style="list-style-type: none"> • Action: Allow/Deny • Protocol TCP/UDP/IP • Source IP • Destination IP • Port • Log

Policy Area	Examples
	<p>Layer 2 ACLs template (to enable organizations to specify the network layer ACL rules):</p> <ul style="list-style-type: none"> • Action: Allow/Deny • Protocol TCP/UDP/IP • Source IP • Destination IP • Port <p>IDS/IPS template(to enable organizations to specify network IDS/IPS signatures):</p> <ul style="list-style-type: none"> • Action detect/prevent • Signatures • Learning period <p>Network isolation (to provide physical and logical separation of network resources)</p> <p>Transport security: TLS, VPN</p>
Storage	<p>The storage rules, constraints, and policies apply to all forms of storage, including block, file, queue, database, and key store. Some examples are:</p> <ul style="list-style-type: none"> • Zones (to provide logical separation of organization storage) • Data classification (to provide organizations with data classifications of <i>public</i> or <i>confidential</i>) • Access/encryption (to allow organizations to request access restrictions or encryption of the data at rest) • Isolation • Location • Retention (to allow organizations to define retention intervals) • Backup and recovery (to allow organizations to specify policy for data backup and recovery)
Compute resources	<p>Affinity Composite application constraints Antivirus Patch management Priority Availability (for example, against DoS attacks) Location and placement Isolation for multi-tenancy Platform (HW/SW) integrity and assurance Scope of policies (where they are applicable)</p>
Security	<p>DPI firewall rules (Relation to Network) allow customers to translate their tier-to-tier firewall rules to the service provider server tier firewall.</p> <ul style="list-style-type: none"> • Action: Allow/Deny • Protocol: TCP, UDP, ICMP, IP • Source IP • Destination IP • Port • Log • Threat management (Relation to Network) • Log management (Relation to Audit) • Legal/regulatory controls • Isolation (Related to Affinity) • Audit

Policy Area	Examples
Alarms	Thresholds Ticketing
Monitoring	Events Alerts Thresholds
Location	Specific location (related to legal/regulatory) Availability

Annex C
(informative)**Change Log**

Version	Date	Description
1.0.0	2010-06-18	Released as DMTF Informational

1780
1781
1782
1783
17841785
1786