



CMDB Federation (CMDBf) Frequently Asked Questions (FAQ) White Paper

Version: 1.0.0

Status: DMTF Informational

Publication Date: 2010-05-10

Document Number: DSP2024

Copyright Notice

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

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

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patent owner or claimant, and shall have no liability or responsibility for costs or losses incurred if a standard is withdrawn or modified after publication, and shall be indemnified and held harmless by any party implementing the standard from any and all claims of infringement by a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

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

Abstract

This white paper describes the key concepts of Configuration Management Database Federation (CMDBf) in a question-and-answer format. Topics include explanations of CMDB, CMDBf, and their components and services, as well as recommendations for implementing them. The questions and answers are organized by role (end user, implementers of repositories and applications to be incorporated in a CMDB federation, and Federating CMDB implementer). The FAQ is intended to provide a simple “on-ramp” for those who want to implement and use the CMDBf specification.

CONTENTS

Abstract	3
1 Introduction	7
2 End Users	7
3 MDR Implementers.....	11
4 Federating CMDB Implementers	14
Bibliography	16

Figures

Figure 1 – Archictectural Components of CMDBf.....	10
--	----

1 Introduction

This white paper is intended for use by people trying to apply the *Configuration Management Database Federation (CMDBf) Specification* ([DSP0252](#)) in real-world business scenarios. The content of this white paper is organized into sets of questions and answers much like Frequently Asked Questions (FAQs).

This white paper is intended to address the following audiences:

- end users, such as IT managers and IT administrators
- MDR (Management Data Repository) implementers
- Federating CMDB implementers

To make it easier to find content of interest to a particular audience, this document is organized by audience type. Questions are not repeated from section to section. If your question is not covered in a particular section, check the other sections.

2 End Users

Q: What is a CMDB?

A: CMDB is an acronym that stands for Configuration Management Database. It is derived from the Information Technology Infrastructure Library (ITIL) best practices. The [ITIL® V3 Glossary](#) defines CMDB as “A database used to store Configuration Records throughout their Lifecycle. The Configuration Management System maintains one or more CMDBs, and each CMDB stores Attributes of CIs, and Relationships with other CIs.” (See [ITIL® V3 Glossary](#), pages not numbered, and [DSP0252](#), clauses 1 and 3.)

Q: What is a CI?

A: CI is an acronym that stands for Configuration Item. The [ITIL® V3 Glossary](#) defines a CI as “Any Component that needs to be managed in order to deliver an IT Service. Information about each CI is recorded in a Configuration Record within the Configuration Management System and is maintained throughout its Lifecycle by Configuration Management. CIs are under the control of Change Management. CIs typically include IT Services, hardware, software, buildings, people, and formal documentation such as Process documentation and SLAs.” (See [ITIL® V3 Glossary](#), pages not numbered.)

Exactly which components and attributes are tracked in a CMDB varies from organization to organization depending on the level of control exercised by configuration management and the needs of the organization. (See [DSP0252](#), clauses 1 and 3.)

Q: What type of data is in a CMDB?

A: ITIL practice recommends that a CMDB focus on the data used to manage the configuration of IT resources that support services, including changes to the configuration of these resources. [ITIL V3, Service Transition](#) (pp. 92-93), suggests six general categories of data:

- Service lifecycle CIs, which include CIs such as business cases, service lifecycle plans and test plans that show a provider’s services, how they will be delivered, benefits expected, costs, and so on
- Service CIs that relate to specific service processes such as capabilities, models, packages, and acceptance criteria

- 39 • Organization CIs, which define organization constraints such as regulatory requirements and
40 business strategies
- 41 • Internal CIs, such as data center assets and software
- 42 • External CIs, such as customer requirements and agreements
- 43 • Interface CIs that are required to deliver service over a service-provider interface

44 In practice, most organizations do not track all these CI categories in the CMDB. Rather, they pick those
45 most relevant to their business. Of these categories, "Internal CIs" usually represent the core of the
46 CMDB.

47 ITIL V3 introduces the concept of a comprehensively integrated system for managing services. Instead of
48 a single database for managing configuration, many federated CMDBs fit together in the Configuration
49 Management System (CMS), Service Asset and Configuration Management (SACM) and the Service
50 Knowledge Management System (SKMS). In CMDBf, a Federated CMDB is designed to support this view
51 of service management as an integrated system. To support the wide range of entities that appear in
52 CMS, SACM, and SKMS, the items that CMDBf refers to are not limited to traditional CIs.

53 In addition to CIs, CMDBs also contain relationships between CIs. A relationship is a link between two CIs
54 that identifies a dependency or connection between them. For example, applications may be linked to the
55 servers they run on. Typically, IT services have many links to all the CIs that contribute to them. (See
56 [DSP0252](#), clauses 1 and 3.)

57 **Q: What is federation?**

58 A: Federation means some or all data from a number of sources can be viewed as coming from a single
59 virtual source. The component databases of a federated database usually maintain independent control
60 of the data they contain.

61 **Q: Why does one federate data? What are the advantages and disadvantages?**

62 A: Usually data is federated to provide a comprehensive view of an object or set of objects from a single
63 access point. In the case of a CMDB, CIs that are desirable to appear in the CMDB are often managed by
64 other applications with their own store of data about the CI. These applications are often called
65 Management Data Repositories (MDRs). One of the goals of the CMDB is to consolidate data from
66 various MDRs into the CMDB to provide a single virtual repository for managing configuration. This is an
67 ideal opportunity for federation because federation allows the MDRs to maintain control of the data on the
68 CIs they manage, but make the data accessible from the CMDB. (See [DSP0252](#), clause 4.)

69 **Q: What are the alternatives to federation?**

70 A: In a limited way, CMDB centralization can be achieved by periodically loading data from the MDRs to
71 the CMDB. This is often called Extract, Transform, and Load (ETL). ETL presents some problems
72 because data stored in multiple places can become stale and out of synch between data transfers.
73 Increasing the frequency of transfers or even initiating a transfer in real time with every change may
74 diminish the stale data problem, but the overhead for transfer can be considerable. ETL systems often
75 transfer data from the MDR that is never accessed from the CMDB. In this case, the system expense
76 incurred to transfer data that is not used is wasteful. In a federated system, the CMDB queries the MDR
77 for data when it is needed by the CMDB. (See [DSP0252](#), clause 4.)

78 **Q: When should federation be used?**

79 A: Because federation avoids stale data and frequent ETL operations, federation is most suitable when
80 the data in the MDRs changes rapidly and the CMDB must have access to the current values. When the
81 MDR contains data on vast numbers of CIs but the probability that the CMDB will need to access any
82 individual CI is low, federation avoids transferring large quantities of data that is never accessed.

83 Nevertheless, loading data may be desirable in some situations. Loading data incurs the overhead of the
84 load into the CMDB, but the overhead for query to the CMDB is reduced because queries are always
85 local. In situations where the data in the MDR remains relatively unchanged over time and the CMDB
86 must access data from the MDR frequently, loading the MDR data may yield better performance with
87 fewer resources. Loaded data will also perform faster when many complex queries involving comparison
88 of data from several MDRs must be executed. A narrow bandwidth network connection between the MDR
89 and the CMDB may favor periodic loads over federation, although this approach usually requires careful
90 analysis of the circumstances. (See [DSP0252](#), clause 4.)

91 **Q: How much data should I federate?**

92 A: The amount of data to federate depends on the services to be supported and the level of configuration
93 control desired. ITIL documentation uses the example of computer keyboards. Most of the time,
94 keyboards are a commodity that is not tracked. But at the United Nations, where the language of
95 keyboards affects the business of the organization, keyboards are tracked.

96 The only guideline for what data should be included is business relevance. Usually, a CMDB contains
97 only CIs that are subject to the change control process. For example, servers in the datacenter almost
98 always appear in the CMDB because they are under change control. If workstations and laptops are not
99 change controlled, they often do not appear in the CMDB. Placing a CI into the CMDB causes additional
100 cost. If there is no return on that cost in the form of more efficient IT services, there is little reason to
101 include the CI in the CMDB.

102 When deciding whether or not to federate an MDR, the same considerations apply. If the data in the MDR
103 is likely to affect IT services, the MDR is a good candidate for federation. (See [DSP0252](#), clause 4.)

104 **Q: What is CMDBf?**

105 A: CMDBf is an acronym that stands for Configuration Management Database Federation. The CMDBf
106 was an industry consortium formed in 2006 to develop a specification for CMDB federation. The
107 consortium was not affiliated with any standards organization. The CMDBf published a specification in
108 October 2007. The CMDBf consortium then donated the specification to the DMTF, and the CMDBf
109 consortium was officially dissolved. The DMTF CMDB Federation Working Group was then formed to
110 shepherd the consortium specification through the DMTF standards acceptance process. The DMTF
111 published the standard as [DSP0252](#) in June 2009. Since the dissolution of the consortium, CMDBf is
112 often used as an abbreviation for “CMDB Federation” and [DSP0252](#) is often referred to as the “CMDBf.”

113 **Q: What problem is the CMDBf trying to solve?**

114 A: The specification ([DSP0252](#)) is meant to support a federated CMDB that provides a single aggregate
115 view of an IT resource, even if the data is from heterogeneous sources. Such a federated CMDB can
116 support many scenarios, including

- 117 • Provide an accurate description of IT inventory from a combination of configuration and asset
118 information
- 119 • Reflect changes to IT resources across diverse repositories and data sources
- 120 • Compare expected configuration versus actual configuration
- 121 • Enable version awareness, such as in the following examples:
 - 122 – Coordinate planned configuration changes
 - 123 – Track change history
- 124 • Relate configuration and asset data to other data and data sources, such as incident, problem,
125 and service-level records. This category includes:

- 126 – Integration of change management and incident management with actual CI status and
- 127 configuration information
- 128 – Analysis of root causes of incidents, problems, and SLA breaches and assessment of
- 129 impacts of planned changes based on dependency analysis of incident, problem, change,
- 130 and configuration records.

131 **Q: Is CMDBf an official standard?**

132 A: CMDBf, *Configuration Management Database (CMDB) Federation Specification* ([DSP0252](#)), is an

133 official standard of the DMTF.

134 **Q: What are the architectural components of CMDBf?**

135 A: The main architectural components of a CMDBf implementation are the Federating CMDB and MDRs,


136 as shown in Figure 1. The MDRs and the Federating CMDB communicate through Registration and

137 Query web services. These services are defined in [DSP0252](#). A Federating CMDB can itself be an MDR,

138 which allows for hierarchical federations of federated CMDBs. In addition to the Federating CMDB and

139 MDR, the architecture describes a Client, an entity that can query MDRs and Federating CMDBs. (See

140 [DSP0252](#), clause 5.)



141

142

Figure 1 – Architectural Components of CMDBf

143 **Q: Is CMDBf an architecture or a piece of infrastructure?**

144 A: CMDBf is a standard that includes an architecture. The software and hardware used in an
145 implementation of the CMDBf standard is part of the enterprise infrastructure. (See [DSP0252](#),
146 "Introduction".)

147 **Q: What is the difference between "Federating CMDB" and CMDBf?**

148 A: In common use, the term CMDBf refers to the DMTF standard, *Configuration Management Database*
149 *(CMDB) Federation Specification* ([DSP0252](#)). A prominent feature of [DSP0252](#) is the Federating CMDB
150 architectural component. The generic term "federating CMDB" without any qualification can refer to a
151 CMDB with federating features that does not necessarily follow [DSP0252](#). (See [DSP0252](#), "Introduction".)

152 **3 MDR Implementers**

153 **Q: What is an MDR?**

154 A: MDR is an acronym that stands for Management Data Repository (see clause 4.6 of [DSP0252](#)). In
155 general, an MDR is a data repository that is external to a CMDB and contains data that would be useful to
156 integrate into a CMDB. The compelling value of CMDBf is that it provides the interfaces for incorporation
157 of MDR data into a CMDB.

158 **Q: Does an MDR need to be implemented using a relational DBMS?**

159 A: No, any data source that is capable of implementing the CMDBf Query Service, and optionally a client
160 that invokes the CMDBf Registration Service, may be used as an MDR. No assumption is made about the
161 kind of technology that implements an MDR. It could be a relational DBMS, an XML file, or a data stream
162 generated from a management application.

163 **Q: What needs to be done to plug my MDR into a CMDBf-compliant CMDB?**

164 A: To plug an MDR into a CMDBf-compliant CMDB, follow these steps:

- 165 1. Identify data to federate.
- 166 2. Decide how a CMDB will learn about the available data in the MDR.
- 167 3. Decide which items and relationships to make available through a Query Service and which to
168 register.
- 169 4. Configure the Federating CMDB and/or MDRs to recognize each other.

170 **Q: How does a Federating CMDB learn about the available data in an MDR?**

171 A: A Federating CMDB learns about an MDR's available data in the following two ways:

- 172 • The MDR informs the Federating CMDB about resources of interest and when they change.
173 This is called "push mode" because the MDR pushes information to the Federating CMDB. The
174 Federating CMDB may also query the MDR for more detailed information about these
175 resources.
- 176 • The Federating CMDB, either periodically or on demand, queries the MDR's data or data that
177 has changed recently. This is called "pull mode" because the data in the MDR is pulled into the
178 Federating CMDB. In a variant of pull mode, the Federating CMDB maintains no data within
179 itself except metadata describing the types of data each MDR contains.

180 **Q: What is the difference between push-mode and pull-mode federation?**

181 A: Push mode registers instances with the Federating CMDB by the MDR invoking the Registration
182 Service. Pull mode allows instances to be “pulled” from the MDR when a query operation is performed.
183 Depending on the pull-mode implementation, the Federating CMDB may periodically poll the MDRs for
184 data of interest or it may wait until a client requests data.

185 **Q: Do I need to implement push-mode or pull-mode federation?**

186 A: You must implement either push-mode or pull-mode federation, or you can implement both.

187 **Q: What is required in push-mode federation?**

188 A: Push-mode federation requires that MDRs invoke the Registration Service exposed by the CMDB. This
189 means that the MDR has a means to actually “push” instances of data to the Federating CMDB.

190 **Q: What is required by pull-mode federation?**

191 A: Pull-mode federation requires that MDRs implement a Query Service. The Federating CMDB
192 determines when to query the MDR data.

193 **Q: How often should an MDR push data to the Federating CMDB?**

194 A: The specification leaves the question to the discretion of the implementation. For example, some
195 MDRs will push data in near real-time updates, like sending a discovery event; other MDRs will implement
196 different policies, such as periodically pushing data at a frequency, such as once per minute, once per
197 hour, etc.

198 **Q: What does it mean for the quality of the federation when an MDR only supports push mode?**

199 A: Using push mode will necessitate that the actual push operations be done at some frequency. If this
200 frequency is occasional, such as once per night, the data in the CMDB may become “stale,” that is, the
201 data might not correspond to what is actually in the corresponding MDR if the MDR changes between
202 push operations. If the pushes occur in near real-time (that is, the CMDB is notified as soon as relevant
203 instance data is added, modified, or removed), then the choice of push versus pull mode has little effect
204 on the data quality. Depending on factors such as the rate of data change and the frequency with which a
205 given MDR is accessed, the resources required to maintain the federation may vary.

206 **Q: What are the advantages and disadvantages of push mode and pull mode?**

207 A: Push-mode federation has the advantage of simplicity and improved query performance (because
208 queries are often partially or completely processed within the CMDB's local data store, rather than being
209 distributed to possibly several MDRs). Push-mode federation has the disadvantage that data may be
210 pushed to the CMDB whether or not it is needed.

211 Pull-mode federation's advantage is that data is transported only when a client queries it or a CMDB
212 wishes to refresh its cache. Pull-mode federation has the disadvantage that queries may take longer to
213 execute, due to the performance cost of gathering data from remote data sources.

214 **Q: What is the CMDBf Registration Service?**

215 A: The Registration Service is an interface implemented by a Federating CMDB and invoked by an MDR
216 that is used to notify the Federating CMDB about resources available for federation. The Registration
217 Service allows instances of data to be “pushed” to a Federating CMDB. After the instances have been
218 added to the CMDB, they can be used to resolve CMDBf queries made to the Federating CMDB. The
219 federating CMDB may also perform some type of reconciliation during registration, but this is not required
220 or defined by the standard.

221 Q: In the Registration Service what gets “pushed” to the Federating CMDB?

222 A: In the Registration Service, a set of items with their associated records is pushed. At a minimum,
223 properties of the items that are being identified should be pushed. However, the standard intentionally
224 neither limits registration to identifying properties nor requires that identifying properties be present. Other
225 properties that are determined to be useful in the Federating CMDB may be part of the content that is
226 registered.

**227 Q: In the Registration Service used by push-mode federation, does the MDR initiate the operation
228 or does the Federating CMDB?**

229 A: In push-mode federation, the MDR initiates the operation by invoking the Federating CMDB's
230 Registration Service.

231 Q: What is an MDR ID, and how can I determine my MDR ID (so that I can register my instances)?

232 A: An MDR ID is a unique identifier, perhaps globally unique, but at least unique among all the connected
233 Federating CMDBs and MDRs. [DSP0252](#) does not specify a means by which an MDR can determine its
234 own ID. As such, this information could be communicated in a number of ways: through a configuration
235 property, through some out-of-band mechanism, through a user interface, or through an API not defined
236 in [DSP0252](#).

237 Q: How will a Federating CMDB identify my MDR's data model?

238 A: The data model of an MDR can be communicated to a Federating CMDB in a number of ways. Clause
239 8 of [DSP0252](#) defines metadata. It is an optional feature of the standard (that is, metadata may be
240 communicated in ways not specified in the standard or not at all in a standard-compliant implementation).

241 Q: What is the CMDBf Query Service?

242 A: The Query Service is an interface implemented both by MDRs and Federating CMDBs. It is used to
243 query the data in the MDR and Federating CMDB. The Query Service includes operations to query
244 instances of data and provide graph queries, allowing for the retrieval of instances that are linked together
245 with relationships.

246 Q: Is the CMDBf Query Service the same as a SQL database query service?

247 A: The Query Service serves a similar role as a SQL database query service, but the CMDBf Query
248 Service does not depend on the use of SQL or an RDBMS, it explicitly provides a way to aggregate data
249 from multiple repositories, and it contains special support for graph queries. GraphQL is the feature
250 that most notably distinguishes CMDBf Query Service from SQL. A GraphQL request describes the
251 items and relationships of interest in the form of a graph. Constraints can be applied to the nodes (items)
252 and edges (relationships). The GraphQL response contains the items and relationships that, through
253 their combination, compose a graph that satisfies the constraints of the query. (See [DSP0252](#), clause 6.)

254 Q: What are the forms of CMDBf queries?

255 A: The two general types of CMDBf queries are property/value and XPath.

256 Q: Does a Query Service require the use of XPath?

257 A: XPath is an optional capability of the CMDBf Standard Query Service. If XPath is used, it should be
258 used according to the guidelines provided in clause 6.5 of [DSP0252](#).

259 Q: Does an MDR use the Registration Service to register itself with the federating CMDB?

260 A: No, the Registration Service is intended to register items and relationships, not MDRs.

261 **Q: Isn't the Registration Service really a "synchronization" service rather than a "federation"**
262 **service?**

263 A: The CMDBf Registration Service is a synchronization service and, consequently, some might argue
264 that the CMDBf is not a pure federation system. However, even a strictly push-mode CMDBf system has
265 a strong element of federation. The primary intention of the Registration Service is to register the *identity*
266 of items or relationships, so that extended information (which may change much more frequently than
267 identity information) may be obtained directly from the MDRs using the Query Service when the
268 information is required. Implementers have the option of building a strict pull-mode CMDBf system without
269 implementing a Registration Service if they wish to avoid synchronization entirely.

270 **4 Federating CMDB Implementers**

271 **Q: What is the difference between a CMDBf-compliant Federating CMDB and my current CMDB?**

272 A: Most CMDBs provide a record of configuration items and relationships in an IT environment.
273 Implementing CMDBf makes a CMDB interoperable with clients and MDRs that implement CMDBf,
274 thereby enabling access to data from sources not managed within your current CMDB.

275 **Q: What do I need to do to make my implementation CMDBf compliant?**

276 A: Implementing all the required parts of the CMDBf standard ([DSP0252](#)) — usually indicated in the
277 standard specification with the word "must" — qualify an implementation as CMDBf compliant.

278 **Q: What is the basic data organization of the CMDBf?**

279 A: CMDBf organizes data using three data wrappers: Item, Relationship, and Record. Item examples
280 include a computer, application, service, building, and incident record. A Relationship is an Item that
281 represents an association between two Items. Each Item or Relationship aggregates any number of
282 Records of any mixture of types. Each Record is a wrapper around data; the wrapped data uses a format
283 selected by the implementation. Clause 5.5 of [DSP0252](#) describes these elements in more detail.

284 **Q: How does a Federating CMDB know which MDRs to federate?**

285 A: The CMDBf Standard does not specify how to determine which MDRs to federate. The two basic styles
286 are to either configure the MDRs to initiate communication with the Federating CMDB's Registration
287 Service (known as push mode), or to configure the Federating CMDB to initiate communication with each
288 MDR's Query Service (known as pull mode). Such information could be specified in a number of ways,
289 including by manually configuring each repository, reading the service addresses from a property in some
290 configuration file, and looking up the information in a service registry. In [DSP0252](#), clause 5.3.2 describes
291 federation modes and clause 8 discusses the use of metadata to describe services.

292 **Q: What is a federated query?**

293 A: A federated query is a query invoked against one repository (the Federating CMDB) that may retrieve
294 and combine data from one or more other repositories (the MDRs).

295 **Q: Must a Federating CMDB implement a federated query?**

296 A: A Federating CMDB must implement the CMDBf Query Service for data contained within its own local
297 repository. It may optionally extend that query to data contained in other MDRs. The Federating CMDB
298 may support a federated query for only a subset of the data available in MDRs.

299 **Q: Can a Federating CMDB be federated by a different Federating CMDB?**

300 A: A Federating CMDB may both federate other MDRs and be federated itself by another Federating
301 CMDB, to which it appears to be an MDR. Two Federated CMDBs may federate each other in a peer-to-
302 peer configuration, or one may federate the other (which in turn may federate others) in a hierarchical
303 configuration.

304 **Q: What is the relationship between roles and services?**

305 A: The active roles are MDR, Federating CMDB, and Client. The services are registration and query.
306 Each role may implement and/or use services. Specifically, an MDR may use the Registration Service
307 and it may implement the Query Service. A Federating CMDB may implement both the Registration
308 Service and the Query Service, and it may use the Query Service (to implement a federated query). A
309 client may use the Query Service. In [DSP0252](#), clause 5.2 describes roles and clause 5.3 describes
310 services, including which services apply to which roles.

311 **Q: How do I uniquely identify my items and relationships?**

312 A: Each Federating CMDB and MDR should have a globally unique identifier or, at the least, an identifier
313 different from any other repository that is part of the same system of federated repositories. Each
314 Federating CMDB and MDR assigns an ID for each item and relationship that is unique within the
315 repository. The combination of the repository ID and the item or relationship ID will be unique within the
316 system of federated repositories. Clause 5.5.2 of [DSP0252](#) describes the formation of IDs for items and
317 relationships, and it also describes record IDs.

318 **Q: How do I avoid collisions in item identification for items originating from different MDRs?**

319 A: Collisions in item identification are avoided by combining the unique repository ID with the local ID that
320 is unique within the repository. Clause 5.5.2 of [DSP0252](#) describes the formation of IDs.

321 **Q: What is the relationship between CMDBf and the DMTF Common Information Model (CIM)?**

322 A: The DMTF CIM does not address the domain of a federated CMDB. Therefore, there is no relationship
323 between the CIM and CMDBf. However, the content of CMDBf-conformant Federating CMDBs and MDRs
324 may be based on the CIM.

325

326

Bibliography

- 327 DMTF DSP0252, *Configuration Management Database (CMDB) Federation Specification 1.0*,
328 http://www.dmtf.org/standards/published_documents/DSP0252_1.0.pdf
- 329 ITIL® V3 Glossary, *Glossary of Terms, Definitions, and Acronyms*, v3.1.24, 11 May 2007,
330 http://www.best-management-practice.com/gempdf/ITIL_Glossary_V3_1_24.pdf
- 331 ITIL® V3 Service Transition, 05 Jun 2007, by Sharon Taylor, Shirley Lacy and Ivor Macfarlane. The
332 Stationery Office. ISBN 9780113310487.
- 333