

The Federated CMDB Vision

A Joint White Paper from BMC, CA, Fujitsu, HP, IBM, and Microsoft

Version 1.0, 25 Jan 2007

Authors

Dale Clark, CA

Pratul Dubish, Microsoft

Mark Johnson, IBM

Vincent Kowalski, BMC

Yannis Labrou, Fujitsu

Stefan Negritoiu, Microsoft

William Vambenepe, HP

Marv Waschke, CA

Van Wiles, BMC

Klaus Wurster, HP

Copyright Notice

© Copyright BMC, CA, Fujitsu, HP, IBM, and Microsoft 2007. All Rights Reserved.

THIS WHITEPAPER IS PROVIDED "AS IS," AND THE AUTHORS MAKE NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT, OR TITLE; THAT THE CONTENTS OF THIS WHITEPAPER ARE SUITABLE FOR ANY PURPOSE; NOR THAT THE IMPLEMENTATION OF SUCH CONTENTS WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.

THE AUTHORS WILL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF OR RELATING TO ANY USE OR DISTRIBUTION OF THIS WHITEPAPER.

The name and trademarks of the Authors may NOT be used in any manner, including advertising or publicity pertaining to this Whitepaper or its contents without specific, written prior permission. Title to copyright in this Whitepaper will at all times remain with the Authors.

No other rights are granted by implication, estoppel or otherwise.

Abstract

This whitepaper presents the vision for a CMDB that is federated from multiple management data repositories, providing the basis for a set of specifications that will standardize how such federation is enabled.

Status

This whitepaper is an initial draft release and is provided for review and evaluation only. It is being published to solicit feedback. A feedback agreement is required before the working group can accept feedback.

At some future date, the contents may be published under another name or under several new specifications, as shall be agreed by the authors and their respective corporations at that time.

Table of Contents

The Federated CMDB Vision	1
Authors.....	1
Copyright Notice.....	2
Abstract.....	2
Status.....	2
Table of Contents.....	3
1. Introduction	4
1.1 Purpose.....	4
1.2 Problem.....	4
1.3 Solution.....	4
2. Requirements	6
2.1 Illustrative Use Case.....	6
2.2 General Requirements.....	6
3. Architecture	7
4. Services and Data Model	8
4.1 Federated CMDB Services.....	8
4.2 MDR Services.....	9
4.3 Data Model.....	10
5. Summary	10
5.1 Specification Phases.....	10
5.2 Expected Adoption.....	10

1. Introduction

1.1 Purpose

ITIL®-based Configuration Management Databases (CMDBs) are emerging as a prominent technology for Enterprise Management Software. The usefulness of these CMDBs is dependent on the quality, reliability and security of the data stored in them. A CMDB often contains data about managed resources like computer systems and application software, process artifacts like incident and change records, and relationships among them. The CMDB serves as a point of integration between IT management processes (as in Figure 1 below). Data from multiple sources needs to be managed directly or by reference in commercial CMDBs. The purpose of this paper is to describe how such data from multiple sources will be federated into a CMDB. The concepts and terms set forth here will provide the basis for industry-wide specifications to be subsequently written and proposed as formal standards.

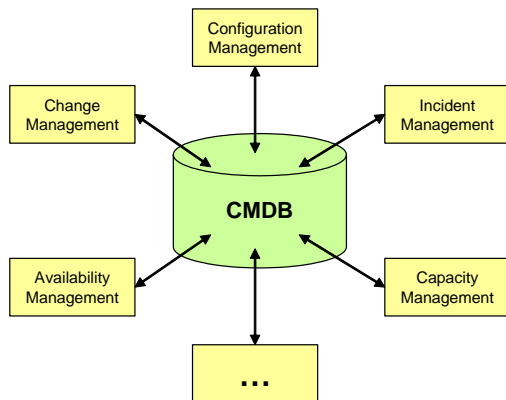


Figure 1 - Role of a CMDB

1.2 Problem

In practice, the goal of federating data is often not met because the various management data are scattered across repositories that are not well integrated or coordinated. There is no standard for providers of Management Data Repositories (MDRs) to plug their data into a federating scheme. This problem exists both for individual vendors trying to integrate with multiple CMDBs, and for customers who need to integrate data from multiple vendors' MDRs.

1.3 Solution

This paper describes the architecture for a CMDB that is created by federating multiple Management Data Repositories. The guiding principles for this architecture are as follows:

1. Enable a variety of data consumers to access a federation of management data through a standard access interface
2. Enable a variety of data providers to participate in a federation of management data through a standard provider interface
3. Provide an approach for reconciling and combining different information about the same resources

This architecture will express the mechanisms for federating the CMDB through standard interfaces without dictating the specific implementation. This will provide the basis for an industry-wide standard that will allow a large number of management data providers to integrate their data into a coherent, seamless CMDB. Consumers of management data will benefit from having a common view of resources and relationships between them, and the ability to use standard queries to retrieve information related to these resources.

From a business standpoint, the existence of such a standard architecture will:

- Drive down the cost of developing adaptors and translators
- Encourage greater participation by MDR implementers
- Improve the quality and integrity of CMDB data
- Reduce the time and effort required to integrate diverse data sources into a common CMDB
- Reduce the customer total cost of ownership of a CMDB by simplifying and centralizing management interfaces
- Enable better IT process management according to industry best practices like COBIT® and ITIL

2. Requirements

2.1 Illustrative Use Case

To make the general problems described in Section 1.2 more concrete we describe a business scenario that is often encountered by IT management personnel as follows.

A problem with one or more resources is causing user transactions to fail or have unsatisfactory response times. The IT Service Desk has created incidents for the response time problems and assigned a subject matter expert to solve these incidents as quickly as possible. The subject matter expert will use resource and relationship data stored in the CMDB, and other management products, to locate the source of the problem. The subject matter expert will determine if the problem resources have had any changes applied recently which may have caused the problem.

IT personnel assigned to such an incident will need to quickly understand what is causing the problem with the application and the impact to business services. This information will help him or her determine how to resolve the service disruption, what priority should be placed on this problem and if there is a need to enlist the assistance of other subject matter experts to assist with the resolution.

The tasks required to solve these incidents would typically be performed by a service desk analyst, subject matter expert(s), and/or an incident manager. Management components (or applications) involved in this scenario include:

- Application (transactions) monitoring application
- Resource monitoring application(s)
- Event analysis application
- Incident management application
- Change management application
- CMDB

The current state of the art in CMDB technology would require all these components to be integrated in proprietary solutions whose reusability would be rather limited.

In the proposed architecture, core information about the applications and transactions (and their supporting IT infrastructure) is available through the Federated CMDB. Related configuration details, performance records, event logs, incidents, and change records are available through the standard queries that each application supports. This architecture provides the basis for an integrated approach to service management.

2.2 General Requirements

In general, solutions to scenarios such as the one described in Section 2.1 share some common requirements. These requirements include:

- Web service interfaces for management data providers to plug into a common CMDB
- Web service interfaces for clients to access such data, including the reconciled common identity of resources, relationships between resources, and related information
- Data aggregation, normalization, federation, and reconciliation
- Secure, authenticated access to these services controlled by an administrative interface

We will consider these interfaces to be the primary focus of our standard architecture, while the implementation of the interfaces and the resulting reconciled resource models will depend on the particular implementation of a Federated CMDB.

3. Architecture

The Federated CMDB architecture will allow IT organizations to create and configure a CMDB from MDRs with various scopes and arranged in various topologies. The external view of the architecture is embodied in Figure 2 below:

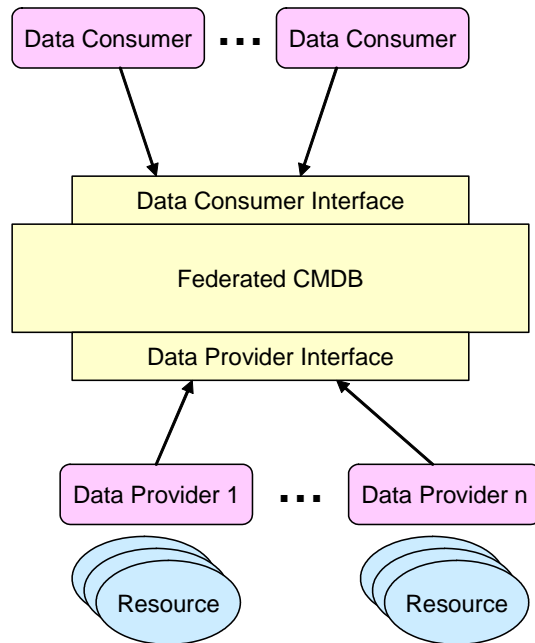


Figure 2 – Federated CMDB High Level Architecture

The elements depicted in Figure 2 are:

- **Federated CMDB**
- **Consumers** who are clients of the data of the CMDB
- **Providers** (MDRs) whose data are federated into the CMDB
- **Data Provider Interfaces** for MDRs to plug their data into the CMDB
- **Data Consumer Interfaces** for clients to access the data from the CMDB

The data consumer interfaces will expose data retrieval operations such as query and subscription. MDRs can plug into the Federated CMDB if they implement the data provider interfaces. These interfaces will allow data sources to register the resources they manage, details about the available data, and how the data can be queried.

4. Services and Data Model

The services to support the above architecture will include the following:

- Federated CMDB Services, which provide for
 - Administration of the Federated CMDB
 - Resource federation, including the registration of resources by MDRs
 - Query of the reconciled resources and their relationships in the Federated CMDB
 - Subscription and notification for both resource and metadata changes in the Federated CMDB
- MDR Services, which provide for
 - Query of the resource records and relationships managed by the MDR
 - Subscription and notification for resource changes in the MDR

A Data Model will represent resource schema (including resource types, identifying properties, and other data properties)

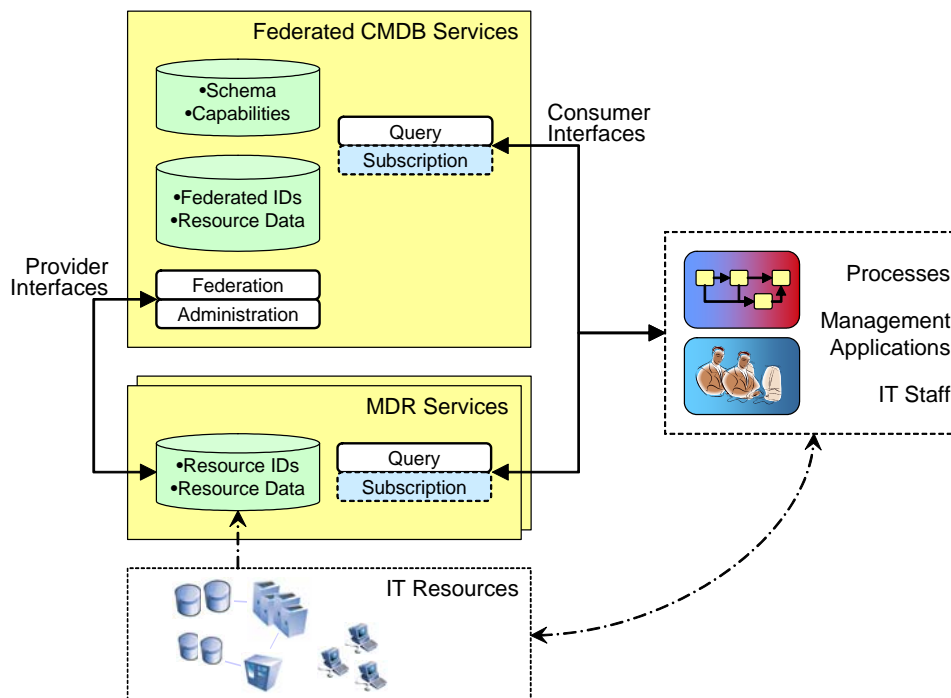


Figure 3 – Federated CMDB Overview

4.1 Federated CMDB Services

There are four Federated CMDB services. Two (federation and administration) are for providers, and two (query and subscription) are for consumers.

- **Administration of the Federated CMDB**

An administration service provides a primary interface of Federated CMDB. All MDRs that participate in the Federated CMDB are registered with the

administration service. Each MDR is associated to schemas (such as resource definition classes) and capabilities (such as support for a query dialect) that it supports. The administration services of the Federated CMDB allow management of the data sources and types that are federated, and control of the content, integrity, and access to the resource data that is registered and federated.

- **Resource Federation and Registration**

A federation service provides the interface to add, modify, and delete resource definitions. The Federated CMDB manages a directory of all registered resources (configuration items, relationships, and process artifacts). It is responsible for reconciling the different perspectives of a resource's identity into a common "resource context". It may also manage additional data properties provided by MDRs when resources are added to the federation service.

- **Resource Query**

A resource query interface allows client applications to retrieve data from the Federated CMDB based on certain criteria. This includes an *instance query* to retrieve items based on identifying properties, and a *relationship query* to enumerate a set of items and relationships based on attributes of the relationships and endpoints.

- **Subscription and Notification¹**

A subscription service provides applications the ability to register for data change events that could occur in the Federated CMDB. Applications needing to respond to data change (update, insert, delete) events will *subscribe* to that event and will be delivered notification when such an event occurs.

4.2 MDR Services

An MDR must register itself to the Federated CMDB administration service. It may also register its capabilities, schemas, resources and relationships. It may support a query service and a subscription service.

- **Resource Query**

A query service may be used to retrieve extended information about one or more registered resources. This includes an *instance query* to retrieve items based on identifying properties, and a *relationship query* to enumerate a set of items and relationships based on attributes of the relationships and endpoints.

- **Subscription and Notification²**

An MDR may support subscription and notification of changes to registered resources. This allows a client of the MDR subscription service to be notified of changes in the MDR data, similar to the Federated CMDB subscription service.

¹ Federated CMDB subscription & notification is not planned for the first phase of the CMDB Federation specifications.

² MDR subscription & notification is not planned for the first phase of the CMDB Federation specifications.

4.3 Data Model

A common data model is essential for successful federation of data from multiple MDRs. A fundamental **Data Model** that underlies the Federated CMDB must include agreed upon definitions of:

- Fundamental Data Types (for example Configuration Items, Relationships, and Process Artifacts)
- Meta Data (information about the structure and format of the data)
- Resource Identity (a consistent mechanism for identifying resources using a resource type and identifying properties, and an associated qualified name or globally unique ID)

5. Summary

5.1 Specification Phases

In order to progress rapidly, the CMDB Federation specifications will be proposed in phases. The first phase will focus on data provider and data consumer interfaces, allowing many management vendors to participate in federation of data, and allowing consumers to query the consolidated resource information from the Federated CMDB and its registered MDRs.

The interfaces in the first phase are primarily intended to support non-federated queries of a single repository, though implementations may support federated queries with these interfaces. Later versions of the specification may define extensions to these interfaces that improve the capabilities and/or performance of federated queries.

Subsequent phases of the CMDB Federation standard may define interfaces for:

- Subscription and notification interfaces for the Federated CMDB and MDR services
- More granular security mechanisms
- Storing information that associates resources with services that can be launched in the context of the resource.

5.2 Expected Adoption

The CMDB Federation specifications will allow management data providers to write a single integration interface that can be used with any compliant Federated CMDB.

Adoption and standardization of these specifications will reduce the cost of ITIL process management and accelerate the adoption of ITIL practices. This will result in improved IT services and customer satisfaction, as the promise of ITIL is realized.