

SIX CRITERIA FOR AN EFFECTIVE CMDB ARCHITECTURE

What do enterprises as diverse as CARFAX, a large retailer, and an infrastructure services provider have in common? Each has adopted IT Infrastructure Library® (ITIL®) service management best practices and implemented automated IT management solutions to achieve important business goals. Those goals include minimizing service disruptions, cutting costs, driving IT efficiencies, and facilitating regulatory compliance.

At the heart of each company's IT management strategy is a configuration management database (CMDB). The CMDB stores all the components of the IT infrastructure as configuration items, maintaining not only detailed data on each configuration item but also data on the relationships among them. The CMDB also maintains management data on the configuration items, including their incident and change histories. By integrating this wealth of data into a central repository, the CMDB offers insight into the cause-effect relationships across management data types.


The CMDB enables integration and automation of processes within and across IT service support, IT operations, and IT asset management. As such,

it provides a strong foundation for Business Service Management (BSM), a comprehensive and unified platform for running IT.

The integration provided by the CMDB smoothes workflow across groups and permits automation of end-to-end processes. This can help IT to improve service quality, manage services more efficiently, and achieve continuous compliance. To provide these high levels of integration and automation, a CMDB must be built on an architecture that meets six important criteria: federation, flexible information model definition, standards compliance, support for built-in policies, autodiscovery, and strong access controls.

1. Federation

By definition, the CMDB provides a single source of truth about the IT environment, and, as such, it is the central repository for data describing the IT infrastructure. But putting all infrastructure information into a single database is just not feasible. There are too many infrastructure types, elements, types of management data, and varying granularity levels per data type.



The CMDB enables integration and automation of processes within and across IT service support, IT operations, and IT asset management.

The practical approach is to unite a collection of CMDBs and other data stores into what ITIL refers to as a Configuration Management System (CMS). In this approach, individual CMDB instances — that is, data stores that specialize in management functions along various boundaries of IT infrastructure and operations management — collectively contribute to an overall enterprise CMDB.

The ability to unite several data stores requires a federated approach that must be designed into the enterprise CMDB architecture and cannot be added as an afterthought. A CMDB built on a federated architecture provides access to a wide range of information without having to move all the data to or replicate it in the CMDB. For this approach to work, the data stores must segment their contributory slices of the enterprise CMDB along well-defined lines. These lines of delineation must satisfy requirements of data exchange, data integrity, and data access.

At a retailer of women's apparel and other items, the CMDB stores fundamental information about the IT environment and provides pointers to other essential and more detailed information in other data stores. Configuration item relationship and management information allows the staff to tie assets to incidents and problems, correlate incidents, and do root-cause analysis. Through federation, the CMDB provides access to information the staff needs to manage asset lifecycles more effectively. This helps to ensure that the company doesn't continue to pay for support and maintenance on assets that have been retired.

A critical requirement for federation is a strong data reconciliation capability to ensure that data aggregated from multiple sources is accurate and consistent. Data reconciliation eliminates duplicates so that there is only one configuration item for each entity. Data reconciliation also ensures that data aggregated from multiple sources are attached to the correct configuration items.

2. Flexible Information Model Definition

Information models for a CMDB can be based on two distinct approaches. The top-down approach takes a desired view of an organization and implements a meta-data model for that view in the CMDB. It follows up by ensuring that all management applications comply with the meta-data model. The bottom-up approach relies on normalizing lower-level data sets to develop a meta-data model.

Most IT organizations opt for the bottom-up approach. That's because it allows existing management data sets to be rolled up into meta-data models easily, reducing implementation effort and speeding time to value. The resulting meta-data models are management-function and application agnostic. Consequently, they are easier to manipulate than actual lower-level data sets that are tied to specific management functions driven by specific applications. The bottom-up approach also has the advantage of easier acceptance because it can be implemented without the disruptive organizational and cultural changes required for top-down approaches. A well-architected CMDB supports both approaches, giving IT the flexibility it needs in implementing a CMDB that best matches its requirements.

3. Standards Compliance

Federation implies multiple CMDB instances and, therefore, multiple data segments. Interoperability among instances and segments is a must. That requires standardized data exchange mechanisms to ensure data integrity, protection, and controlled access. Consequently, the CMDB architecture must support open standards, such as XML, and Service Oriented Architecture around Web services. Standards support enables interoperability with disparate data stores while ensuring that data sets do not violate the overall integrity of the meta-data definitions that the IT organization has developed for its enterprise CMDB.

4. Support for Built-In Policies

A well-architected CMDB maintains policies that document the standards used to build, update, enforce, and track continuous compliance for a given service and its supporting components. A service might encompass an application, middleware, system utilities, databases, network devices, and an operating system. Supporting infrastructure components might include Web servers, database servers, application servers, network devices, and clients.

Standards must include details on data sets, such as configuration, installation, performance, and uptime. Policies may be dynamic and may change depending on factors such as time, number of users, and service level agreements (SLAs).


A well-architected CMDB also maintains process models. Because the IT environment typically changes over time, these process models must be dynamic and adapt automatically to those changes.

Because it maintains policies and process models, the CMDB plays a pivotal role in policy-based process automation. This automation greatly speeds process execution while enforcing the use of best-practice processes. For example, an infrastructure services provider that is focused on card-based payment transaction services, electronic payment systems, and international financial information, credits its CMDB with enabling IT to execute all active and proactive processes — such as release, change, and SLA management — at extremely high-quality levels in a very short timeframe and with exceptional performance.

5. Autodiscovery

The CMDB must automatically discover all the assets in the IT infrastructure, details about them, their physical and logical relationships to each other, and their relationships to the services they support. Federation supports autodiscovery by serving as a means to gain detailed insight into any given infrastructure component.

CARFAX, the leading provider of vehicle history information, relies on automated tools to discover configuration items in the IT environment and feed them into the CMDB. The tools also capture the logical dependency mapping of components, identifying which IT components comprise enterprise applications.



The CMDB must automatically discover all the assets in the IT infrastructure, details about them, their physical and logical relationships to each other, and their relationships to the services they support.

Traditionally, IT has used autodiscovery to rapidly populate inventory information. The latest generation of autodiscovery solutions also periodically scans the IT environment to provide live configuration insight into a given component at different points of time. Live discovery combined with time-series snapshots are vital to any analysis performed on components and the services they support.

6. Strong Access Controls

In the IT world, unauthorized or undesired access and changes can result in service disruptions or downtime. Consequently, security and access controls play an essential role in CMDB design and implementation. Access policies can be used to create profiles and access controls for users and groups.

A CMDB must conform to security standards to prevent the implementation of any unauthorized changes to its data set. These standards can be documented through directories to identify who has access to which data sets and who is authorized to perform actions on which data sets. A CMDB, with this built-in, role-based access control, supports user authentication through directories to claim access rights.

Foundation for Configuration Automation

A well-architected CMDB can serve as a strong foundation for IT organizations to gain visibility, accountability, and control over their services infrastructure. When combined with discovery and decision support applications, the CMDB provides the foundation for a holistic platform that automates the configuration management of services while ensuring continuous compliance with corporate policies, government regulations, and industry standards and best practices.

For more information about the BMC Atrium CMDB, visit www.bmc.com/atrium.

BUSINESS RUNS ON I.T. I.T. RUNS ON BMC SOFTWARE

Business thrives when IT runs smarter, faster, and stronger. That's why the most demanding IT organizations in the world rely on BMC Software across both distributed and mainframe environments. Recognized as the leader in Business Service Management, BMC offers a comprehensive approach and unified platform that helps IT organizations cut cost, reduce risk, and drive business profit. For the four fiscal quarters ended Sept. 30, 2009, BMC revenue was approximately \$1.88 billion. Visit www.bmc.com for more information.