Deploying the CMDB for Change & Configuration Management

MAY 2007

CA and BearingPoint

Table of Contents

Executive Summary

Challenge

Businesses are placing high demands on IT to adapt to changing conditions. However, IT organizations tasked with providing increased service levels at lower costs do not have the resources available to reinvent themselves with every technological or regulatory change. Without frameworks in place to leverage automation and best practices, these IT organizations are consumed with the day-to-day operations of IT — with little time and few resources left to develop new services that add value to the business.

Opportunity

Adopting a best practice framework, such as the IT Infrastructure Library (ITIL®), is critical for IT departments looking to address internal and external pressures, control costs and enable business advantage. By providing data to support and enhance the ITIL processes, the Configuration Management Database (CMDB) plays a critical role within the ITIL framework. The CMDB is a repository of Configuration Items (CIs), which are elements within the IT infrastructure including hardware, software, business applications and services, and the relationships between them. A mature CMDB enhances all ITIL processes with particular emphasis on the Change and Configuration Management processes.

Benefits

A fully deployed CMDB managed by the Change and Configuration Management processes offers key benefits including:

- Reduced downtime due to root-cause analysis of service outages

- Minimized operational risk due to change impact analysis

- Improved contingency planning due to visibility of IT components underlying business services

- Improved service integrity due to monitoring of IT components related to critical services

- IT process efficiency and cost-effectiveness as a result of sharing a consistent set of IT service, application and infrastructure data

Pressures on IT Departments

CIOs are pressed to deliver and demonstrate business value to the organizations they support. They must maintain the quality, efficiency and productivity of day-to-day IT operations despite decreasing IT resources, budgets and time. More than ever, businesses are placing high demands on IT departments to adapt to changing conditions. At the same time, new regulations and technological innovations introduce complications and waylay IT organizations tasked to provide increased service levels at lower costs.

Without a framework in place to leverage scale, automation and best practices, the resulting operating structure may be rigid and reactionary, increasing IT costs by forcing IT to reinvent itself with every technological or regulatory change. CIOs may find themselves focused solely on the day-to-day operations of IT — with little time and few resources left to develop new IT services that can enable business advantage.

Adopting Best Practices Can Help Address the Pressures

A popular and proven response to the pressures on IT is the adoption of a service-oriented culture that focuses on outcomes and customer experience. Process improvement based on the ITIL best practice framework is the bedrock of this cultural shift. Best practice frameworks provide wide-ranging benefits of standardization, consistency in delivery, guidelines and improved quality of service. In short, adopting a best practice framework of IT processes is one of the key strategies an IT department can undertake to address internal and external pressures, control costs and add value to the business.

CMDB's Role in Change and Configuration Management

The CMDB plays a critical role within the ITIL framework by providing critical data to support and enhance Service Support and Service Delivery processes. The CMDB is a repository of CIs, which are elements within the IT infrastructure including hardware, software, business applications and services, and their relationships to each other. While a mature CMDB increases the effectiveness of all ITIL processes, it is most closely associated with and usually deployed within the context of the Change Management and Configuration Management processes.

In order for the CMDB to be deployed, the Change and Configuration Management processes must be addressed together. The Change Management process strives to minimize or prevent service disruptions with standardized methods and procedures for the efficient handling of all Requests-For-Change (RFC). While Change Management handles the workflows of change requests, analysis, approvals, scheduling, implementation and reviews, it is dependent on Configuration Management for a clear, accurate picture of the IT environment including specific Configuration Items and their relationships to each other.

The goal of Configuration Management is to provide comprehensive and accurate records of CIs throughout their lifecycles, including identification, control, maintenance and verification. When implemented as part of a Change and Configuration Management initiative, the CMDB maintains accurate details of IT components and the services they support and enables the IT organization to deliver cost-effective, high-quality IT services.

Configuration Management activities related to the CMDB include:

- Creating a list of key attributes (e.g. physical and logical characteristics such as the number of processors, the amount of RAM installed, size of hard drive, relationships to key IT and business services, relationships to other components, etc.) to collect and maintain for each CI

- Defining the relationships between individual CIs, the relationships between CIs and IT services and the relationships between IT services and business services

- Tracking the current status and history of each CI

- Verifying and ensuring that the CI attributes are complete and accurate

- Maintaining an authorized state for comparison to actual state

Configuration Management is distinct from IT Asset Management (ITAM) in that in addition to tracking physical assets such as routers, servers and ports, Configuration Management includes IT services and functions, and highlights their relationships and interdependencies. Although some of this information must be gathered manually, relationship mapping tools can automate the discovery of relationships for both physical IT components and logical CIs like business services. Often, CMDB projects focus initially on ITAM discovery tools so that the CMDB can be populated with an initial list of discovered CIs. In these cases, attention to relationships between the discovered CIs, or between the discovered CIs and applications and services, comes later.

When Change and Configuration Management processes are integrated and fully functioning, IT can truly understand and control the IT infrastructure. Incidents can be prioritized so that resolution is focused on those with true business impact. Proposed RFCs to a CI can be assessed and approved with a clearer picture of the CI's relationships to other devices and impact on IT and business services. Unauthorized changes to the infrastructure can be minimized and their resulting costly business service outages avoided. IT services can be defined, measured and ensured more precisely. The capacity of IT resources and services can be more accurately anticipated and planned for. To achieve these benefits, Configuration Management maintains the current status of the dynamic IT environment in the CMDB, with Change Management maintaining data integrity with updates to the CMDB.

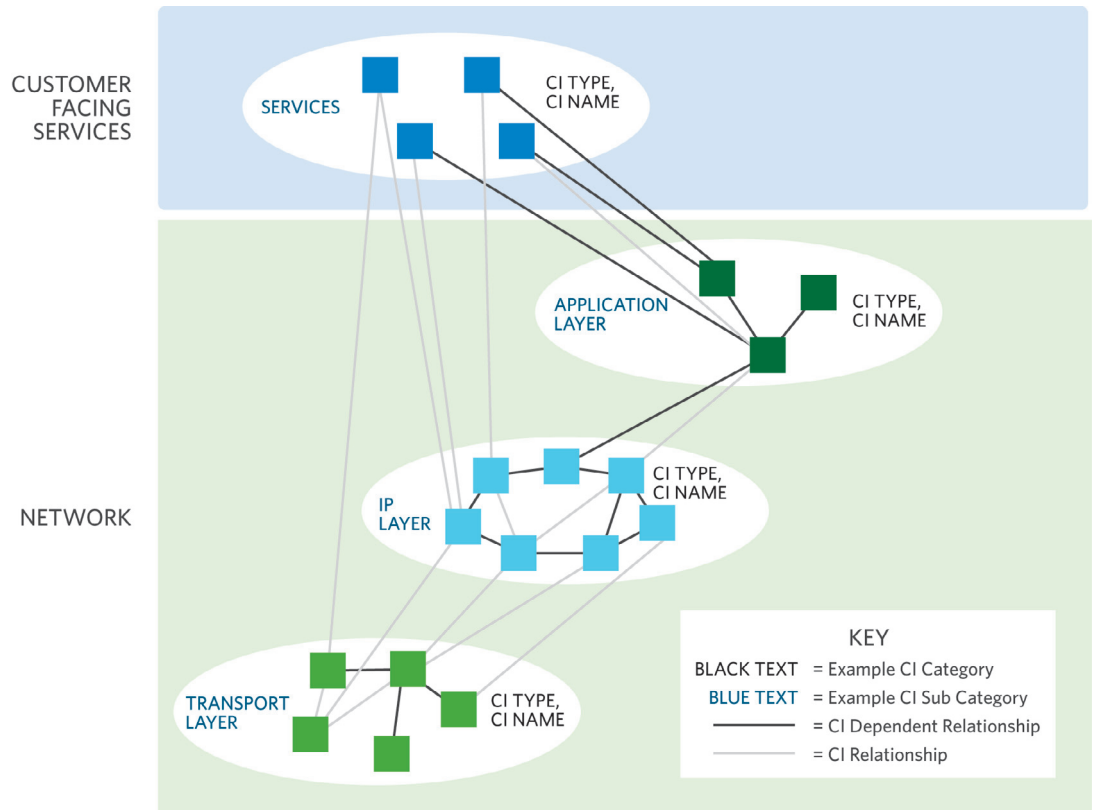# Benefits of Change and Configuration Management

A fully deployed CMDB (See Figure A) managed by the Change and Configuration Management processes offers important benefits including:

- Reduced mean time to repair (MTTR) through service outage root cause analysis
- Increased mean time between failures (MTBF) through better change planning and more accurate impact analysis
- Adherence to contractual obligations
- Minimized operational risk through improved change impact analysis
- IT accountability to the user community
- Heightened security by controlling CIs in use
- Improved contingency planning due to a thorough understanding of the CIs underlying critical business services
- Improved Infrastructure integrity by identifying, controlling and correcting CI exceptions (instances where the discovered CIs differ from previously known CIs)
- Control of assets through verification of CI records against the discovered infrastructure
- Gain in IT process efficiency and cost-effectiveness through sharing of a consistent, single and accurate set of CIs across the entire infrastructure
- Improved planning as the status of the entire infrastructure is kept current. No unplanned changes are made and all changes to CIs (i.e. new incidents, software deliveries and patch updates) will be recorded in the CMDB
- More accurate risk assessments as the relationships and dependencies between CIs are known
- Improved root-cause diagnosis and resolution through a graphical view that provides visibility to identify upstream and downstream CIs impacted and related to the problem

CIs can range from configurable service views to nodes on the network identified by IP addresses. CIs each have a unique name, belong to a category and type and have defined relationships with other CIs.

CUSTOMER FACING SERVICES

SERVICES

CI TYPE, CI NAME

APPLICATION LAYER

CI TYPE, CI NAME

NETWORK

IP LAYER

CI TYPE, CI NAME

TRANSPORT LAYER

CI TYPE, CI NAME

KEY

BLACK TEXT = Example CI Category
BLUE TEXT = Example CI Sub Category
―――― = CI Dependent Relationship
―――― = CI Relationship

CMDB Capabilities

The CMDB is the software that tracks CIs, their attributes, their relationships and their interdependencies. It is the centralized repository for critical configuration data within the organization and the foundation for reaping the benefits of many IT Service Support and Delivery processes. Without a CMDB, the opportunity to fully exploit Change Management, Configuration Management and other ITIL processes is missed. Without Change and Configuration Management, the CMDB is likely to remain a simple IT asset repository or a set of manual inventory snapshots.

The architecture of a CMDB, including
Visualization and the CI population
and maintenance disciplines of
Synchronization, Reconciliation,
Federation and Relationship Mapping.

CMDB ARCHITECTURE
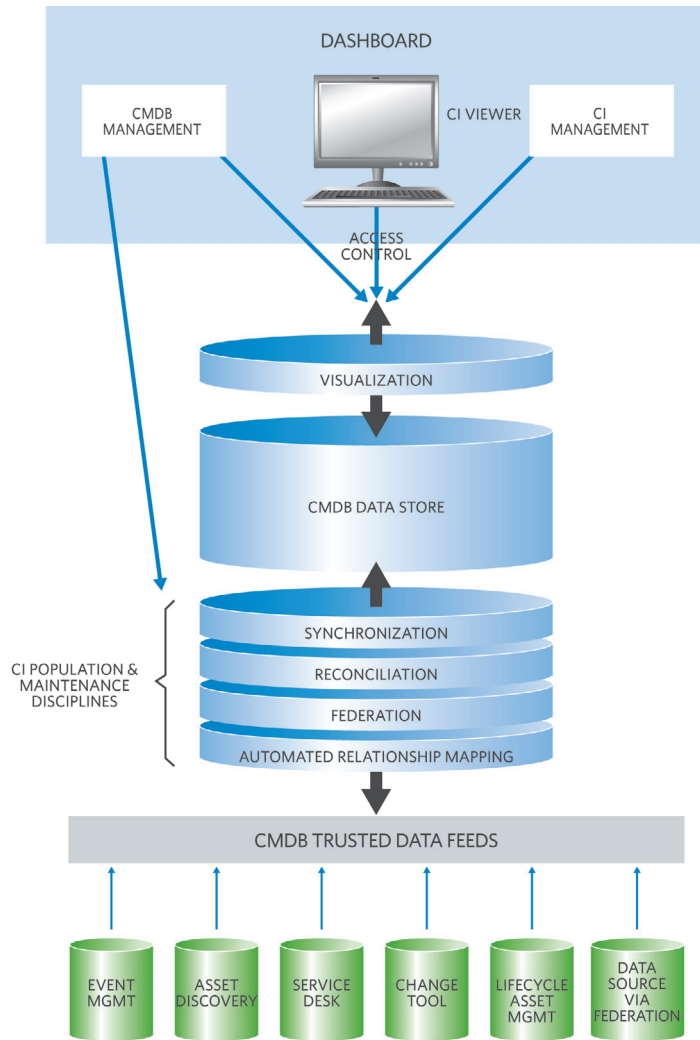
A full-featured CMDB includes the following fundamental capabilities (See Figure B):

- **Federation** Integrating multiple configuration data sources by duplicating and/or connecting to (and then maintaining) data from a source to the CMDB. This is accomplished using Extract, Transform and Load (ETL) techniques. Alternatively, the CMDB may view remote data without duplicating it from its original source. In this case, the CMDB should be able to create relationships between both resident and remote CIs. All data sources should be "trusted sources" — sources where data integrity, cleanliness and validity is known. Infrastructure and application component discovery tools, service desk and network and system event management tools are likely sources. Other sources include contract management systems and service level management systems. Ultimately, the list of data sources is dictated by the problems the CMDB is enlisted to help solve, and the IT processes the CMDB supports.
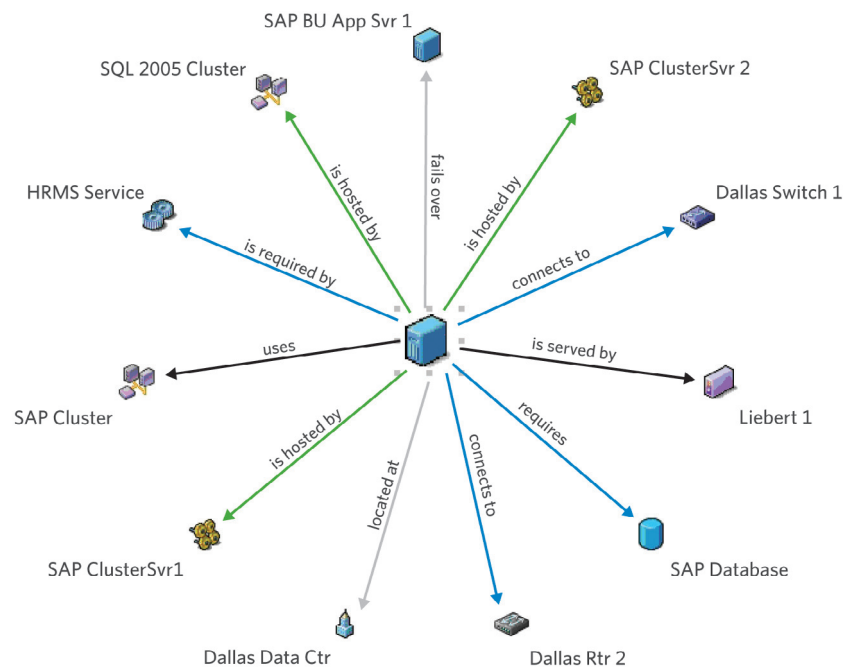
- **Reconciliation** Recognizing and then rationalizing data from the same instance of a CI that originates from multiple external sources. Attribute data for a single CI can often come from multiple sources, both when the CI is initially populated into the CMDB and when updates are made over time. Thus, the CMDB is most effective when it is able to consolidate and reconcile information from multiple sources and IT domains. This ensures that all IT services are working from the same consistent set of configuration data. Reconciliation of large volumes of CIs can consume many resources and therefore should be automated as much as possible.

- **Visualization** Illustrating the peer-to-peer and hierarchical relationships between business services, IT services and CIs using well-designed, clear graphical and reporting interfaces (see Figure C). A CMDB offering visualization must have a sophisticated ability to record, maintain and display the topology and relationships of the physical and logical IT services, applications, servers, network and end user devices (e.g., PCs, laptops, PDAs).

FIGURE C

A CMDB should be able to record, maintain and visualize relationships between CIs.

CMDB VISUALIZATION

SAP BU App Svr 1

SQL 2005 Cluster

SAP ClusterSvr 2

HRMS Service

Dallas Switch 1

is hosted by

fails over

is hosted by

is required by

connects to

uses

is served by

SAP Cluster

Liebert 1

is hosted by

located at

connects to

requires

SAP ClusterSvr1

SAP Database

Dallas Data Ctr

Dallas Rtr 2

- **Synchronization** Ensuring the CMDB is in step with infrastructure monitoring and tracking tools. For example if a Network and Systems Management system detects the outage of an application fail-over server, the appropriate CI in the CMDB can be updated to reflect the outage. The impact to the business service can be determined through a visual display and contingency measures initiated.

- **Relationship Mapping** Mapping relationships and dependencies from IT services (or business applications) down to the CIs that enable those services and populating that data in the CMDB. Relationship and dependency mapping tools should recognize the logical relationships between mapped CIs as well as any changes to those relationships.

CMDB Implementation Considerations

The implementation of a CMDB must be thorough and the CMDB consistently maintained. If, for example, CIs quickly become out of date or incomplete, then decisions based on the CMDB are likely to be wrong, damaging the credibility of the CMDB and of the IT department. The following list should be considered by organizations looking to leverage a CMDB to resolve specific business issues:

- Ensure the CMDB tool is developed as part of the Change and Configuration Management process initiative — do not assume the CMDB is the process. The project's focus is Configuration Management of which a component is the CMDB. Similarly, ensure the project has sufficient sponsorship as multiple data owners will be required to participate in the CMDB population if it is to be successful. The value proposition for the CMDB must be clearly articulated to these stakeholders.

- Consider a phased introduction of CIs starting with a pilot project. Focus initially on the physical CIs (e.g. servers, routers) that are accurately and consistently discovered, and then incorporate the higher level business applications and services starting with the most critical. Add these to the CMDB through collaboration between IT and Business process owners to ensure that the CIs are clearly linked to solving targeted business issues. Succeed with this undertaking and then expand.

- Determine and understand required configuration data. Planning for the CMDB implementation will require agreement and definition of the CI categories, types and the actual CIs that should be captured and stored in the CMDB.

- Document existing data sources and their accuracy. Many CMDB efforts get bogged down in the initial integration with, and migration of, data from existing repositories that are often inaccurate, incomplete and even useless. Identify and prioritize the repositories that require the least amount of data cleansing and provide the biggest bang for the buck when leveraged by the CMDB.

- Determine the approach to federation. That is, what data should be fed into the CMDB, through a one time migration or regular updates, and what data will remain in the domain-specific tools. Related to this will be a decision on whether a staging area for data loading, CI reconciliation and cleansing is required. Update cycles are often different for each IT technology domain. Therefore, reconciliation and CMDB CI refreshes will have to be sequenced and timed to ensure that the data being updated in the CMDB is as close to "real-time" as possible.

- Prepare for reconciliation by rationalizing and cleansing the data to be reconciled. Remove as many inconsistencies as possible. For example, settle on a name for a manufacturer and stick with it. Don't use an abbreviation for the manufacturer name on one CI record and the full name in the next. Inconsistencies like this will cause automated reconciliation to fail. With clean and consistent data, automated reconciliation will save a great deal of manual effort.

- Pilot the project to ensure accuracy of all data feeds and to verify reconciliation logic. Once the data is verified as accurate, the project can progress to the next step where the data is reconciled, tested and then considered "trusted" and eligible for inclusion in the CMDB. By illustrating the ability of the CMDB pilot through carefully selected Use Cases, for example providing a list of all impacted components due to a router failure, the credibility of the CMDB and its ability to support root cause analysis is proven.

Roles and Responsibilities of Change and Configuration Management Providers

- The duties of the Configuration Manager include maintaining a centralized CMDB. Training for this responsibility is essential. CMDB maintenance activities include: manual CI reconciliation, CI to IT service mappings, creation of new CI classes and relationships, developing reports and maintaining federated data sources. Data source owners should be regarded as stakeholders and be responsible for the push or pull of data from their respective technology domains.

- The duties of the Change Manager include filtering, accepting and classifying RFCs, obtaining the required authorization for a change, planning and coordinating the implementation of a change, reviewing implemented changes to ensure they meet their objectives and convening Change Advisory Board meetings to review and/or approve urgent RFCs. Training for this responsibility is essential.

SECTION 4

Conclusions

Deploying the CMDB can help organizations improve their Change and Configuration Management processes. The upside to a well-planned CMDB that fully exploits the principles of federation, reconciliation, synchronization and relationship mapping is considerable. Enhanced root-cause and impact analysis leads to improved service levels and increased user and customer satisfaction. As change and configuration processes evolve and mature, change itself becomes more controlled and predictable, resulting in fewer outages caused by unauthorized or poorly planned changes. Improved traceability meets compliance objectives. Ultimately, the IT infrastructure becomes more agile because services are developed, tuned and realigned quickly without jeopardizing on-going business activities.

Despite the temptation to quickly achieve these considerable benefits, a CMDB implementation should not be rushed. It must be carefully planned, starting with an analysis of current process and data sources paired with a pilot deployment, in collaboration with the owners of the data that will populate the CMDB. Time invested in addressing federation, reconciliation, synchronization and relationship mapping issues before deployment is time well spent. The result is a fully functional CMDB, which confirms an IT organization is moving up the maturity curve — and marks that organization as a world class IT service provider.

About the Authors

Adrian Oldfield
BearingPoint

Adrian Oldfield is a Manager in BearingPoint's IT Strategy & Transformation London practice. He has been an IT professional for 18 years, the last five of which have seen him specialize in IT service management technologies and practices, including the design and development of configuration management databases. His wider IT experience spans application development through to program management and management consultancy.

Robert Sterbens
CA, Business Service Optimization

Robert Sterbens is Director of Product Marketing for IT Service Management (ITSM) solutions at CA, in the Business Service Optimization business unit. He is responsible for developing marketing strategies, business plans and managing industry analyst relations. Robert is currently a member of the itSMF USA Advisory Board where he provides technology and thought leadership in the field of ITSM.

Marv Waschke
CA, Business Service Optimization

Marv Waschke is VP, Development and Senior Technology Strategist, CA, in the Business Service Optimization business unit, where he managed development of CA's service desk product. He was a representative to Network Management Forum trouble ticketing standards committee. For CA, Marv chaired the DMTF Support Work Group, and now sits on the Service Management Language Working Group and the CMDB Federation Working Group. Marv's opinions on IT service management and standards appear regularly in his blog 'Iterating on IT Service' found at http://ca.com/us/blogs

CA, one of the world's largest information technology (IT) management software companies, unifies and simplifies complex IT management across the enterprise for greater business results. With our Enterprise IT Management vision, solutions and expertise, we help customers effectively govern, manage and secure IT.

Transforming
IT Management