





Driving a New IT Reality

THE POTENTIAL OF THE CMDB

A federated CMDB provides a single source of critical information that enables IT to better support business goals. The right data in the right format, reconciliation capabilities, and processes to ensure quality data are the keys to making the CMDB a successful reality in your IT organization.

By JONATHAN MARKWORTH

There's a lot of excitement today about the potential of a configuration management database (CMDB) — and, in particular, a federated CMDB — to help organizations drive IT efficiency and truly benefit from their investment in IT assets. But there's also a lot of confusion, misinformation, and concern about what a CMDB is and how much business value it can actually deliver. Some common questions include:

- > What is a CMDB, anyway? And why is federation so important? Is the CMDB a myth, reality, requirement, or just another sales and marketing ploy?
- > I already have dozens of data repositories, management consoles, and inventory databases spread throughout my organization. Can I control all the systems of record I'll have to deal with in a federated model? Can I implement a federated enterprise CMDB without bringing my organization to its knees?
- > My people seem to be drowning in information, but they're unable to find out what they really need to know. How will adding another tool to the mix help my staff manage the business more proactively?

The relationships between CIs in the CMDB help unravel the complex interdependencies in the IT infrastructure and link them to business services.

- > Implementing a CMDB and maintaining configuration item (CI), attribute, and relationship data can't be easy. Will I end up with a CMDB that propels the business forward or one that just consumes already scarce resources without a return on investment? Will the CMDB become shelfware or a powerful tool for my business?
- > Is this just going to be one more thing that my overworked staff has to deal with?

WHAT IS A CMDB?

Several definitions are being offered, but it is generally accepted that the CMDB is a data repository relating to the IT infrastructure, and contains information that identifies each CI, describes attributes about each CI, and denotes the relationship between CIs. Your organization may be planning to simply deploy a CMDB as a centrally available repository of auto discovery information collected by electronic inventory tools. However, the CMDB really should be designed as a solution that enables IT to support critical business functions. As simple as this sounds, the design and implementation of the CMDB should be tailored to the goals of your organization.

In the context of critical business functions that are enabled by various IT services, CI information stored in an isolated data repository is virtually useless. Centralizing the repository is a good first step in deploying a CMDB. However, CMDB deployments must also have relationship information that links stored data back to the context of the business. Only then will the CMDB have the maximum effect on any anticipated process and performance improvements. Relationships are required to fully realize the benefits of a CMDB across your IT organization. It's the relationships between CIs in the CMDB that help unravel the complex interdependencies in the IT infrastructure and link them to business services. And it's the relationships that unite management applications and processes in a cohesive and efficient manner to drive higher levels of quality, predictability, and service impact assessment across your enterprise.

FEDERATED CMDB — MYTH, REALITY, REQUIREMENT, OR PLOY

Wouldn't it be great to have a single, all-knowing, all-powerful, and self-maintaining tool to manage every facet of information about the IT infrastructure? Wouldn't it be great to start over fresh — to replace your



legacy enterprise management infrastructure with tools that did everything within a fully integrated platform? The reality is that most organizations have inherited (or deployed) dozens of applications, tools, utilities, data stores, hardware platforms, and management frameworks that perform one or more of the service management functions. Each of these has a data repository that provides information to support some critical function within your current environment. In the context of the CMDB, the collection of the data repositories related to these tools contains most of the critical CI attributes that can be used to build relationship data. The key question is then how do you leverage the investments you've already made to build an integrated data repository such as a CMDB?

One method is to export the CI identification, attribute, and relationship data from each of these data sources and combine it into a single database. Over an extended period of time, this method can be very difficult to maintain and only becomes more complicated as the number of data sources grows.

An alternative approach is to build a CMDB that consolidates identification, attribute, and relationship data, making it readily available

to all the IT processes that need it and without requiring all of the data to be copied centrally. In this federated model, the CMDB holds configuration items and a critical subset of data, but then links to other sources of related data such as service desk tickets, service

The federated model can facilitate the introduction of an enterprise CMDB that spans the whole IT organization.

level agreement definitions, or even management consoles. Deployed correctly, the federated model can facilitate the introduction of an enterprise CMDB that can span the whole IT organization, allow for a phased transition plan from legacy systems — if that is even necessary — and preserve the ability of the IT organization to continue day-to-day operations with fewer interruptions.

TOO MANY SCREENS, TOO MUCH DATA, TOO GENERIC A FORMAT

The volume and scope of the information stored in a CMDB can become overwhelming if not carefully managed. Providing the right information, at the right time, to the right people, and in a format that encourages the

correct human or technical response is critical. Think of the user interface of the CMDB in terms of a modern fighter plane. Many years go into modeling the interaction between the pilot and the cockpit displays so that in a high-stress situation the pilot can find exactly the information he or she needs.

Deploying a CMDB without an effective and efficient reconciliation engine places the entire program at risk.

When a virus attack hits or a major component fails, processing information from the CMDB, including the extended CMDB data, isn't much different.

The technicians responding to incidents or implementing changes are the pilots of your IT environment. Give them too many screens or too much data in too generic a format and their decision-making ability and productivity will decline, and the probability of making a mistake will increase. The CMDB is no longer useful for them. The fear of crashing and burning will keep them using their old, familiar spreadsheets and databases, no matter how complete and accurate your CMDB is advertised to be. The CMDB must support views in which information is efficiently tailored for its intended use.

THE ABSOLUTE IMPORTANCE OF RULES-BASED RECONCILIATION

The quality of data will vary from one source to another, and sources may have conflicting or overlapping data that use different labels for the same data. Resolving the conflicts and reconciling the data is one of the most important and most challenging aspects of maintaining the quality of a federated CMDB. In many cases, the volume of data and the number of sources make manual reconciliation impossible. So a robust, rules-based reconciliation process is a must. This is an extremely important requirement. Deploying a CMDB without an effective and efficient reconciliation engine places the entire program at risk.

PROCESS IS KEY TO MAINTAINING DATA QUALITY

When you're using a CMDB to better support critical business functions based on configuration data, the quality of that data is crucial. Consequently, you have to define ways to measure quality. You must also consider data and data management process quality over time. An effective tool that can be used is the RollThroughputYield formula:

Bad data * Bad data * Perfect data = Less than perfect data

It doesn't matter that the last step in a process produces perfect data if the steps feeding into it are of poor quality. It is critical that you are resolute and aggressive in measuring and continuously improving the processes that produce and maintain data,

5 KEY BENEFITS OF A FEDERATED CMDB

1 Achievement of a centralized view without the cost of moving all data to a single, monolithic data repository.

2 Ability to leverage existing investments that currently isolate critical information.

3 Visibility into the complex interdependencies throughout the IT infrastructure.

4 Greater efficiency, leading to higher productivity and lower IT costs.


5 Enhanced ability to comply with government regulations that impact IT.

so that you can ensure the quality of your data. You must continue to improve quality by looking at every process to identify and eliminate redundant or error-producing steps. Lack of a structured process improvement methodology is most definitely the number one cause of poor CMDB quality that eventually results in implementation failure. Perception truly is reality in this case. The audience that receives information from your CMDB probably won't care too much about why the data looks wrong — just that it's wrong.

THE FUTURE — BINDING DISPARATE PROCESSES TOGETHER

Implementing a CMDB is complex, but the benefits are significant. A federated CMDB built on an adaptable technology enables you to start binding disparate IT systems, data sources, and processes into a cohesive environment that facilitates business goals. Having a single source of critical information accelerates the speed at which you can execute, because it provides every person in the IT organization with the information they need to work efficiently and effectively — from the service desk agent answering the phone or the technician refreshing a router, to the mid-level manager working on next year's budget or the CIO developing long-term IT strategies. ●

ABOUT THE AUTHOR



Jonathan Markworth is a managing consultant with CompuCom Systems, Inc. (www.CompuCom.com), headquartered in Dallas, Texas. Jonathan has successfully performed a variety of functions in the IT industry during the last 20 years. He is certified in ITIL and Six Sigma, and he is a certified Project Management Professional (PMP). As part of CompuCom's Integrated Infrastructure Management practice, Jonathan and the CompuCom Center of Excellence team are responsible for deploying effective solutions for change, configuration, release, and asset management.