bit2bit americas DEVICE42

The CMDB is Dead, Long Live the CMDB!

Confronting IT information management requirements in the new digital pipeline

EXECUTIVE SUMMARY

The Configuration Management Database (CMDB) concept started to provide IT organizations the information they need to manage their resources and applications effectively. It's defined by ITIL 4 as a tool that "is used to store configuration records throughout their lifecycle and... maintain the relationships between [them]."

It's a useful idea, yet many CMDBs have failed to live up their billing, leaving many organizations disenchanted. This doesn't change the fact that IT teams will always need a place to get a comprehensive and detailed view of what makes up their enterprise.

The real problem isn't in the concept of a CMDB itself. It's about how to build and maintain this kind of resource, no matter what you call it. IT teams typically have legacy systems along with growing numbers of cloud systems--not to mention shadow IT resources. These organizations find it difficult to locate all their IT assets to initially populate the CMDB while constant changes in IT assets render most CMDBs virtually impossible to keep up to date. IT teams also struggle to track the diverse relationships between IT assets, although this information is required to successfully migrate, improve, or decommission these assets without suffering service interruptions.

This white paper details the challenges modern organizations face in creating and keeping up a CMDB. It then explains how IT teams can get the complete, timely asset information they need to maintain whatever CMDB they decide to use to always have accurate, up-to-date date on their IT estate without endless hours of manual work.

The CMDB: Origins, Purpose, and Challenges

The concept of a CMDB first arose in the 1980s as part of the IT Infrastructure Library (ITIL), a best practice framework developed by the government of the United Kingdom and other organizations to help manage and develop controls for IT services. ITIL principles and standards called for creating and maintaining a database to track and manage IT assets and services—the CMDB.

A CMDB is meant to store all the configuration data about IT assets (referred to as configuration items (CIs)) throughout their lifecycle. These assets can include hardware, software, systems, facilities, and sometimes personnel. Configuration data can include data on class and attributes of each CI (such as type, owner, and importance), a history of changes to the CIs, relationships, interdependencies between CIs, and much more.

A typical CMDB stores information about thousands of CIs. It's not uncommon for organizations greater than 1,000 users to have tens of thousands of CIs. Companies with a large IT footprint can have 100,000 or more CIs. The goal of the CMDB is to provide an

organization with the information it needs to make better business decisions and run efficient IT service management (ITSM) processes within such complexity.

Yet Many CMDBs Fail to Deliver

The CMDB is often positioned as a perfect resource to solve all an organization's asset management challenges, but, industry statistics tell us that only 25% of organizations get meaningful value from their CMDB. The problem has little to do with what CMDB is in use or how the data it may contain is stored. The real problem is in how to get the most complete, accurate data on Cis and their relationships to start with, and how to keep that information up to date after that.

Difficult Discovery

MDB tools offer great features for storing and retrieving data about assets, but they don't offer a clear path of how to get the data required. That's up to other technologies, including sheer human effort. At first, that wasn't an insurmountable problem. When CMDBs were first introduced, organizations primarily had on-premises data centers. IT teams could walk into a server room and quickly create a manual inventory of servers and other assets.

Now, most IT enterprises resources are the result of years of adding, changing, an/or removing IT assets. Many organizations have tens of thousands of servers in scattered locations. These are often a heterogeneous mix of legacy systems as well as cloud-based resources and shadow IT systems that individuals or teams have commissioned without oversight from IT leadership.

The sheer volume and rate of change that has occurred and continue to happen within these systems make it impossible to track CIs effectively using manual inventory methods. Some clues may exist in spreadsheets or in multiple tools, but most IT teams have IT asset documentation that is far from complete. As a result, many organizations suffer with CMDBs that incomplete and not fit for purpose.

The Problem with Updates

Iven if organizations do manage to create a somewhat detailed CMDB through a combination of tools and manual effort, data about assets in large organizations change constantly. As soon as the CMDB is populated, it's out of date.

Inability Visualize Dependencies

Most organizations currently do not have tools that enable them to easily see and evaluate the complex relationships between CIs in the CMDB, even though knowledge about these dependencies is necessary to optimize or migrate systems to the cloud. Visualizations can help organizations see clusters of systems that are rarely used. With that knowledge, organizations can decommission unused assets to reduce IT operating costs as well as improve security by reducing the IT attack surface. But such intuitive visualizations of complex IT are hard to find.

What is Needed

No matter what CMDB organizations use to store the details about their IT estate, it's clear that they need to start with more complete information and keep it complete as they move through time. The real issue isn't the CMDB itself, it's how to get and maintain the data in it. Behind CMDB technologies, organizations need automated, thorough, and reliable IT infrastructure and application discovery technology together that includes clear dependency mapping for all the Cis discovered.

Comprehensive Auto Discovery and Continuous Updates

An automated discovery solution can identify CIs and their related data across an organization. It can then auto-populate any CMDB, even if the IT team using it originally has limited knowledge of what they actually have in the IT estate. It should include RESTful APIs to enable your organization to customize the process of data-loading and mapping. For assets that are not connected to the network, and therefore cannot be discovered automatically, the solution should make it easy to import data about CIs from third-party sources, too.

Visualizations Show Relationships Between CIs

Visualization tools should be available within the tools that support the CMDB to enable an organization to better understand the relationships between Cis quickly. Highly customizable reports and inquiry screens can enable organizations to visualize computer rooms and racks with power and capacity heatmaps, understand their full IP connectivity and power chain, as well as see hardware, software, and application dependencies. The best discovery tools feature dependency mapping that shows logical dependencies and connections with limits on noise from too much complexity. They do this by allowing for the creation of logical affinity groups as shown in Figure 1.

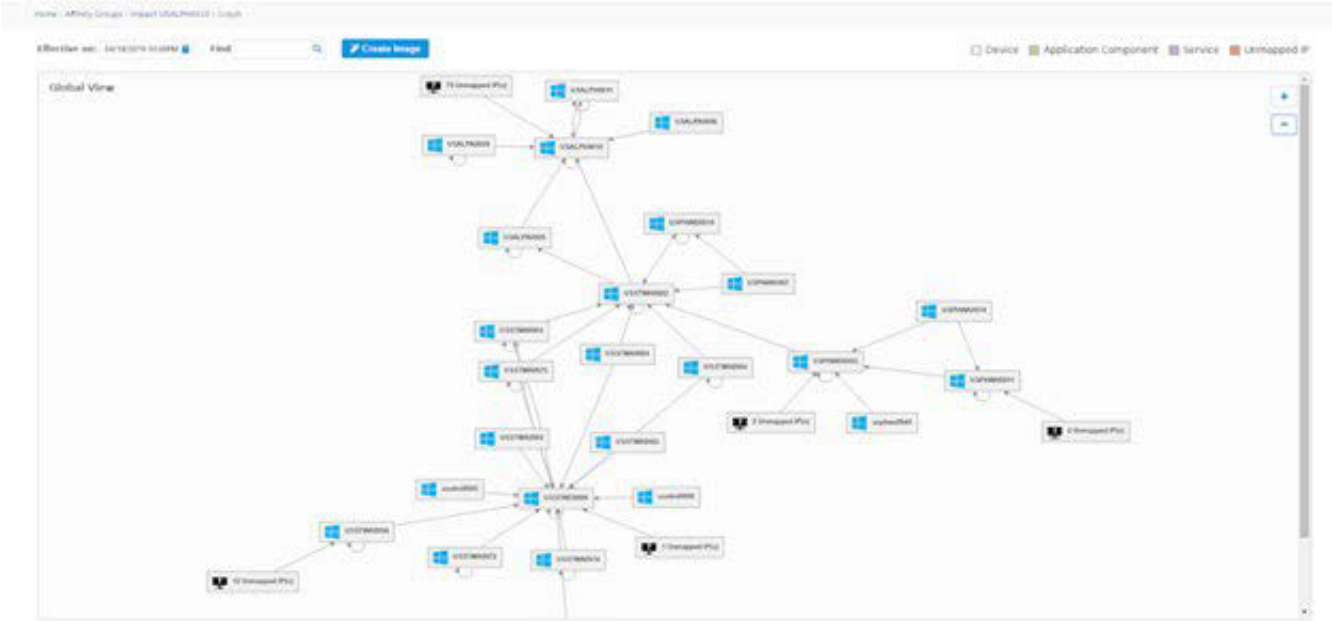

Figure 1: An affinity group shows dependencies in a simplified, but useful way with just enough detail.

Accelerate with Discovery and Dependency Mapping Designed to Support the CMDB

Many CMDB products require organizations to define many fields for the Cis they include, which can make it a challenge to provide complete information simply by manual means. The best discovery and dependency mapping solutions include pre-built and customizable fields that discover and add in CI details immediately through the automated discovery process. The automated process defines users, groups, and passwords directly from the UI and define object-level permissions of Cis for granular control. From there, all the IT team must do is export the complete data to the CMDB tool of their choice.

Critical CMDB Use Cases

If it's fed the right data, a CMDB can be a useful solution for a wide range of use cases across an organization. For example,

- Create a central point of reference for all their IT assets and infrastructures that makes it quick and easy to answer questions such as: How many of a certain type of hardware does the organization have deployed? How many installations of Software X currently exist within the organization? How many software licenses does the organization need to purchase to cover all its usage?

- Reduce risks of datacenter modernization efforts by providing a comprehensive accounting of IT resources and infrastructure that need to be modernized and eliminating the potential for forgotten assets.

• Quickly locate capacity for load response scaling in the face of constantly changing demand.

• Better manage all aspects of data center operations. A CMDB acts as the central repository for infrastructure information and serves as the one source for managing daily operations including client support, trouble ticket resolution, trouble ticket planning, and more.

The CMDB Requires Automated Discovery and Dependency Mapping

The CMDB is a solid concept, but it offers only a partial solution at best. Often, in practice, the CMDB is like having a great library with only some of the books partially documented in it. To be the single source of truth required for the modern enterprise, a CMDB requires a comprehensive discovery and dependency mapping solution working to support it with the data it needs to be the roadmap to the future it is promised to be. Without the right discovery and dependency mapping solution to feed the right data into the CMDB, even those who follow ITIL best practices face significant risk simply because the data within their CMDB is flawed or incomplete.

Avoid Penalties

Automated discovery and dependency mapping solutions enable organizations keep track of legal, contractual, and regulatory obligations. For example, modern discovery tools can track active software licenses within the IT estate to make sure that organizations aren't going beyond their license limitations. With the right information discovered and kept in the CMDB, audits are no longer an issue.

Increase Efficiency and Lower Costs

Discovery and dependency mapping platforms provide visibility into what hardware and software their employees use—or don't—to support wise decisions based on actual operational needs. This data helps organizations prevent unnecessary purchases or determine which assets can be decommissioned, but it's almost impossible to get without the right technology.

Improve Security

What IT teams don't know can cause big problems. Without the proper dependency mapping and discovery platform, organizations rely on manual documentation and searching for the assets they have. Often times, resources such as servers, network devices and applications simply remain invisible in place in large IT companies, long after they are no longer useful. These redundant assets not only increase cost, but each increases the potential surface area vulnerable to attack in the enterprise. With a modern discovery solution, organizations gain visibility into all of their resources, including the never-before-seen.

Boost Compliance

When regulators arrive, an organization can't be sure it's compliant if it doesn't know what assets it has. Having a CDMB is a great start, but, unless that CMDB is kept up to date, the audit is usually accompanied by a huge manual effort to see what's really there. A modern discovery and dependency mapping platform takes this guesswork out of the picture to show how data is used by resources throughout the organization and how it's kept safe.

Conclusion

The need for a single source of truth across the complex, hybrid IT estate has never been greater. IT is at the heart of all business activity—from the daily churn of financial markets to the production, logistics and distribution that make food available. However, after decades and layers of IT development, growth, and innovation, most organizations struggle with business technology that isn't fully documented or even fully visible to the teams that manage it.

• That's a big problem because it's impossible to protect or improve what you can't see.

• The solution to the CMDB conundrum isn't in the CMDB itself—it's in how it starts and how it is maintained. To overcome these challenges, organizations need automated discovery and dependency mapping technology that...

- Automatically discovers diverse configuration items (CIs) across a hybrid enterprise and provides details on each CI.
- Connects to and updates CMDB resources automatically every time it runs.
- Enables clear and simplified visualizations of complex dependencies with enough detail to be useful, but not so much to be confusing.

The best solutions provide this level of visibility and integration data in a way that's easily brought directly into IT Service Management (ITSM) and reporting visualization tools. The right solution automates and simplifies the process of starting and maintaining an accurate and useable CMDB.

About Device42

Device42 is the most comprehensive agentless discovery and dependency mapping platform for hybrid IT available today. Device42 can continuously discover, map, and provide insight to optimize infrastructure and applications across data centers and cloud, with accurate views of your IT ecosystem. Once discovered, Device42 intelligently groups workloads by application affinities, dramatically reducing the effort required to create move groups, capturing all communications.

Customers in more than 60 countries including Global 2000 clients and systems integrators use these capabilities as they manage and modernize their IT infrastructure and application landscapes and adopt DevOps practices.

Start your free trial today and see what Device42 can do for you.