

DATABASE

TRENDS AND APPLICATIONS

Solutions for the Information Project Team • www.dbta.com

Volume 21, Number 5 • May 2007

The Growing Importance of the Configuration Management Database

By Glenn O'Donnell

The configuration management database (CMDB) has become one of the hottest topics in IT operations. A growing number of organizations are looking to better manage IT resources through a federated collection of these databases, which hold information about files, applications, servers and other systems as well as their relationships to one another.

Several business and technical factors are coming together to create the interest in a federated CMDB. As more companies face increased competition and closer scrutiny of their information assets as a result of regulatory pressures, there is a growing need for more disciplined, efficient management of IT services and a closer alignment of those services with the actual information needs of the business. Up until recently, artificial barriers between front- and back-office operations existed because organizations supported key business processes with a few monolithic application environments that did not easily share information. These barriers made it difficult for IT to provide the information business users needed - especially as their requirements changed.

Today's business processes are increasingly interconnected, frequently crossing functional boundaries. People inside and outside the company must be able to share and act on a common set of information to remain productive. So, IT is being asked to get rid of the information silos the monolithic application environments had created and provide integrated information that supports the new cross-functional business processes.

As IT works towards meeting these requirements by building a more integrated information infrastructure, they are undertaking initiatives such as data center consolidation and automation and adopting ITIL best practices, in order to improve the efficiency and discipline of IT operations. All of these

initiatives require accurate configuration information to be effective. So, it is understandable that the CMDB, the core element of the configuration management process, has become such an important issue facing IT organizations.

The Role of ITIL

Organizations are demanding the function provided by a CMDB because it is essential to new technology initiatives such as service-oriented architectures (SOAs) and ITIL implementations. SOAs are providing the application environment necessary to support the tight interconnection of business processes across functions. Many organizations are also leveraging the ITIL libraries, which recommend a set of best-practice processes to address key areas of IT service management, to help enforce discipline and governance on an SOA.

While the ITIL libraries reference the CMDB extensively, technical specifications are purposely avoided in order to maintain neutrality. The ITIL libraries do state that the CMDB should "hold the relationships between all system components, including incidents, problems, known errors, changes, and releases. It should also contain information about employees, locations, suppliers and business units."

Even though there is no schema or model defined for the CMDB by ITIL, there are enough use cases defined for each process in the ITIL libraries to build the CMDB properly. And, emerging standards like Distributed Management Task Force (DMTF)'s Common Information Model (CIM), an object-oriented model for unifying and extending existing standards such as SNMP, are theoretically interoperable and do illustrate what detail is needed.

Federated Object Model

It is important to note the differences

between a CMDB and a federated CMDB, and highlight the significantly higher value that a federated approach provides. A federated CMDB is really not a database but a federated object model that links all trusted sources of configuration information. These sources are often existing standalone CMDBs that cover network, system, application, or other infrastructure components, as well as users, policies, geography, relationships and references. Because it is likely that many management tools are already in place, the federated CMDB should leverage these tools and databases as the trusted sources. Organizations that take advantage of existing CMDBs and tools will not only get more accurate configuration data, but also a higher return on investment and much quicker time to value.

Of course, there will inevitably be conflicts between sources so the federated CMDB must also enforce reconciliation policies to help determine which source is the trusted source and resolve the conflicts consistently. And integration issues among the sources, an age-old plague, will continue to require attention in order to overcome the challenges to automate configuration management processes.

Organizations must be careful of any approach that claims federation but relies on a central database of replicated information from these sources. This approach is not a federated model and could lead to issues or errors related to the currency of the data due to deficiencies in the replication process.

Defining Relationships

The real power of the federated CMDB resides in the relationships between the configuration items rather than the configuration items themselves. The CMDB must define how the configuration items come together to produce and consume information services. The question for many companies is how

much information to put into the CMDB regarding the relationships. Some CMDB efforts get bogged down by attempting to describe the data itself rather than sufficient detail to uniquely describe the source and enough context to define the relationships at a high level. The relationship map is more than just the physical links. Anything that provides context is potentially useful.

The abstracted model created by these relationships is critical to the success of the CMDB. It will become the engine to drive automated discovery of configuration information, for example. The relationships impart meaning to applications and services that are constructed from the specific underlying components. These abstracted models are more meaningful to the IT and business stakeholders than the individual components. IT must involve the business users in the creation of the model to ensure its relevance and accurately reflect the business perspective.

Automation is a Key

Organizations that relied on manual processes to populate a CMDB often found that the configuration information was outdated even before it was entered. With configuration information at the center of every IT process, it is essential that the CMDB data is complete, accurate and current - otherwise all processes are potentially flawed. With the complexity of today's average IT environment and the amount of change that occurs daily, it is simply not practical to manually reconcile, normalize, and populate configuration information when configuration changes are made. The only way to do this effectively is to automate the processes.

Companies should take note of the availability of new auto-discovery software. The tools help organizations automatically populate and maintain configuration information. Outside of the standard instrumentation provided for network components through SNMP, there are numerous but less mature methods and technologies for collecting the details of infrastructure components. Programmers can use command-line interfaces like WMI, SSH and Cisco IOS, to generate the details of non-SNMP components like servers, applications and storage. Relying on only one of these mechanisms to provide automation, however, would likely

lead to incomplete data. Therefore, organizations should look to incorporate multiple sources.

While a truly federated CMDB is not easy, it is really the only choice that will ensure success. And automation is central to managing the intricate relationships that make the federation work.

Service Automation

Organizations will find that the CMDB by itself is not very useful, but the process-oriented usage of the CMDB suggested by the ITIL libraries will provide a profound power to transform IT operations. The intelligent application of policy-based automation is the key to many of the operational benefits that are driven from the CMDB. Companies will realize vast improvements in the efficiency and accuracy of IT services through automation.

Several software technologies, including software analysis tools that monitor the status of all configuration components, can help organizations drive the automated enforcement of policies - from service level agreements to IT governance and regulatory compliance policies. The tools will also provide them with the flexibility to decide how automation is carried out, including which policies are completely automated and which require a mix of automation and human approval gates.

Companies should combine these analysis tools, auto-discovery technologies and policy administration tools with the CMDB to ensure that the CMDB is closely linked to disciplined change management procedures. This will allow IT to model and test configuration changes before implementation, which should improve the accuracy of change procedures. IT can also use auto-discovery and policy administration tools together to automatically monitor configuration changes and ensure that unauthorized requests cannot be made. Without the tight link to effective change management provided by these tools, companies will receive very little value from the CMDB.

Conclusion

Organizations that are looking to successfully implement a CMDB should keep the following guidelines in mind. Given the increasing importance configuration information

holds for effective IT service management, it is reasonable for organizations to look at creating a separate position within IT that is accountable for the CMDB. Given the cross-functional nature of the role, the position must be given enough authority to make decisions and initiate action.

ITIL should continue to guide organizations in the development of best-practices IT service management but these organizations should know where the ambiguities exist and be able to resolve them in a way that results in a CMDB that properly supports business process requirements.

Companies should use a federated model for the CMDB but understand clearly what "federated" means. This means that they should avoid any approach that populates a central database by replicating data from disparate configuration sources. Some replication may be unavoidable at first, but highly-distributed, yet tightly linked federation must be an unwavering principle of the CMDB architecture.

Organizations looking for the key to the federated object model should use the technologies available to automate the CMDB as much as possible, creating a strong link to change management in the process.

Without the involvement of the business users, the models that represent the relationships among the infrastructure components and the overall configuration will likely have little relevance. The business users' perspective must be reflected in the CMDB in order to realize the full potential value.

Most importantly, companies must be realistic about where the technology is at this stage and patient about achieving success. Solutions will involve multiple tools from multiple vendors and integration challenges are certain. Pick a specific process which will utilize the CMDB, such as service level assurance, and implement that. Don't just implement the CMDB for its own sake.

By taking the time to fully understand how the CMDB leverages current management tools, organizations will be in a position to improve IT service management, realize substantial efficiency gains and better support the requirements of the business.

Glenn O'Donnell is principal product marketing manager with EMC's Resource Management Software Group.