



itSM Solutions® DITY™ Newsletter Reprint

This is a reprint of an itSM Solutions® DITY™ Newsletter. Our members receive our weekly DITY Newsletter, and have access to practical and often entertaining articles in our archives. DITY is the newsletter for IT professionals who want a workable, practical guide to implementing ITIL best practices -- without the hype.

become a member

(It's Free. Visit <http://www.itmsolutions.com/newsletters/DITY.htm>)

Publisher

itSM Solutions™ LLC
31 South Talbert Blvd #295
Lexington, NC 27292
Phone (336) 510-2885
Fax (336) 798-6296

Find us on the web at: <http://www.itmsolutions.com>.

To report errors please send a note to the editor, Hank Marquis at hank.marquis@itmsolutions.com

For information on obtaining copies of this guide contact: sales@itmsolutions.com

Copyright © 2006 Nichols-Kuhn Group. ITIL Glossaries © Crown Copyright Office of Government Commerce. Reproduced with the permission of the Controller of HMSO and the Office of Government Commerce.

Notice of Rights / Restricted Rights Legend

All rights reserved. Reproduction or transmittal of this guide or any portion thereof by any means whatsoever without prior written permission of the Publisher is prohibited. All itSM Solutions products are licensed in accordance with the terms and conditions of the itSM Solutions Partner License. No title or ownership of this guide, any portion thereof, or its contents is transferred, and any use of the guide or any portion thereof beyond the terms of the previously mentioned license, without written authorization of the Publisher, is prohibited.

Notice of Liability

This guide is distributed "As Is," without warranty of any kind, either express or implied, respecting the content of this guide, including but not limited to implied warranties for the guide's quality, performance, merchantability, or fitness for any particular purpose. Neither the authors, nor itSM Solutions LLC, its dealers or distributors shall be liable with respect to any liability, loss or damage caused or alleged to have been caused directly or indirectly by the contents of this guide.

Trademarks

itSM Solutions is a trademark of itSM Solutions LLC. Do IT Yourself™ and DITY™ are trademarks of Nichols-Kuhn Group. ITIL® is a Registered Trade Mark, and a Registered Community Trade Mark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office, and is used here by itSM Solutions LLC under license from and with the permission of OGC (Trade Mark License No. 0002). IT Infrastructure Library® is a Registered Trade Mark of the Office of Government Commerce and is used here by itSM Solutions LLC under license from and with the permission of OGC (Trade Mark License No. 0002). Other product names mentioned in this guide may be trademarks or registered trademarks of their respective companies.

[Company](#)[Methodology](#)[Products](#)[Contact](#)[Resources](#)

DITY™ Newsletter
Vol. 3.4 • January 24, 2007

Enterprise CMDB

[digg](#) (discuss or comment) on this article. Show your support for DITY!

[Print](#) page



[Subscribe](#) Its Free!



[RSS](#) Stay up to date!



[PDF](#) Pass it around!

A CMDB is a nebulous thing, and can exist in the minds of the organization, 3x5 cards, or any other medium. However, an enterprise CMDB is different. For large organizations only software will do, but not traditional relational database software, what we need is a new breed of software products...and they might be coming...

By [Hank Marquis](#)



I usually tell smaller and less mature IT organizations that they don't need to invest in expensive and complicated CMDB software. I have used Configuration Management principles of knowledge management to help different IT departments using basic office tools like Excel.

A CMDB is a system for managing knowledge more than it is a product, and it is more of an index than a database. This means as long as the amount of data is small, you can get the benefits of a CMDB any number of ways—paper, Visio diagrams, Access databases, peoples heads, etc. And this is exactly what you see in most organizations.

However an enterprise IT organization spans large geographies with hundreds or thousands of

users and locations, and the sheer number of configuration variations requires a specialized CMDB solution. For the record, I don't think such a product is available yet. The requirements lay just outside of today's technology and protocols.

But this is changing. First to change was the realization that a CMDB is not simply an IT Asset Manager (ITAM) "on steroids," but rather that a true enterprise CMDB requires some unique features. Chief among this is that a true enterprise CMDB solution has to dust of an old database concept that is a bit out of style—dimensional modeling. An enterprise CMDB solution has to be based on a dimensional database, a relational database simply can't do the job.

Following I explain what a dimensional database is, and what an enterprise-class CMDB software solution requires: Federation, Reconciliation, Synchronization, and Modeling.

Many vendors are starting to talk this talk, but few understand what it really means, and fewer still deliver products that walk the talk.

Relational vs. Dimensional

Data stored in a relational database is easy to lookup when you know in advance what you will want to see. The common "row and column" approach of relational databases are ideal for On Line Transaction Processing (OLTP).

But a CMDB does not store most of its data; it references data stored externally in other, perhaps relational, databases. And a CMDB is used to provide contextual awareness over non-obviously related bits of data. For example, a common inquiry posed to configuration management might be: "How many users, in sales, use SAP during the last week of the month?"

This kind of query is not well suited to a centralized relational database with pre-built SQL queries. There are just too many possible combinations of data. This type of query has to pull information from many systems, and the data it needs is probably not all nicely lined up in rows and columns ready to query. No, an enterprise CMDB has to be dimensional—a technology that represents data as different dimensions or plains.

The dimensions of a CMDB often include locale (e.g., city, state, floor, etc.), work group like sales, marketing and so on, IT service like SAP or Email, date ranges, and others. Instead of an Excel spreadsheet with rows and columns, think about a Rubik's cube and you begin to get the idea. The logic required is not new, its been around for years in the form of On Line Analytical Processing (OLAP).

Don't get too excited—simply having OLAP does not give you a CMDB for a couple of very special reasons. First, most CMDB data resides outside of the CMDB system. In order to pull this data from many sources requires federation—a new CMDB buzzword you will begin to hear about more and more.

By way of an example of federation and what it requires, lets consider an IT service for project management. Composing this service are human resource information residing in SAP, project management data in Microsoft Project Server, IT asset information in CA Unicenter, and networking hardware resource data discovered and stored in CiscoWorks.

Federation

The first major requirement for dimensional modeling is federation, or referencing data from several sources instead of replicating it. The CMDB is a meta-database, that is, it is a database that references other databases. The issue that drove federation the first time around was data

validity. If you make a copy of something, then what is definitive? The original or the copy? And how to you know if the copy is the same as the original?

The issues around federation are how to connect to heterogeneous data sources, resolve which bits of data are definitive, and then create and store keys with unique data not found in any external data source but still required. For example, data not found in any of these systems might be the name of the IT service and which workers use it.

There also has to be a method to store awareness of the types of data to be found in each federated data source in order to process ad hoc queries.

As you can see, the idea of federation is easy to state as "connecting to multiple data sources", but having a CMDB system that can actually federate is a very tall order indeed. And federation is just one of four equally complicated CMDB technical requirements. Consider this: what if two data stores reference the same data? Which data store is definitive then, and more importantly, how would you determine which is definitive? This is the issue of reconciliation.

Reconciliation

Aside from the issues of simply connecting to heterogeneous and possibly competitive data sources, the big problem with federation for a CMDB is data confrontation. During creation and maintenance of the contextual information in the CMDB meta database key bits of data transfer from federated data sources to the CMDB data store. Since it is common to have multiple applications and systems that overlap and monitor the same IT assets or store similar data, possible data inconsistencies and redundancies arise.

For example, in our sample project management service, perhaps CA and CiscoWorks both store hardware data: Unicenter may refer to a router by name, perhaps "CISCO01". CiscoWorks may be aware of the same router, but not by the name "CISCO01", but rather by its IP address "128.10.0.1"—this is a real problem since there are not two routers but one. How does the CMDB system discover that "CISCO01" and "128.10.0.1" refer to the same single router? Further, how would the CMDB system know its a router at all? This is the domain of reconciliation.

Reconciliation implies adjusting data derived from more than one source to eliminate duplicates and maintain consistency of data. Federation is useless with reconciliation. Adding even more complexity to a CMDB system is the need to handle any changes arising from successful reconciliation, and this leads to synchronization.

Synchronization

Most data stored in federated CMDB systems changes—sometimes slowly, sometimes swiftly. For example, considering our example project management IT service, the name of the project manager (e.g., the "user") might change, or the type of router hardware might change from a Cisco 2504 to a Cisco 7502. Reconciliation has to be able to can resolve these differences to maintain the integrity of the CMDB. But this is IT, not "simple" data warehousing. Nothing in a CMDB should change with a Request for Change (RFC). Thus, a CMDB system that can successfully federate and reconcile data must also be able to alert when it detects unauthorized changes.

This makes the CMDB system require an awareness of approved changes. Then, when the reconciliation engine detects and resolves a change in infrastructure or data, it has to compare this change to an expected list of approved changes and generate an alert if the change is unapproved (e.g., not planned.) This alert brings CMDB data to the attention of its administrators, who need

help visualizing, mapping, and displaying data—modeling.

Modeling

Modeling is mapping and visualizing synthesized relationships that are IT service definitions. Modeling is more than reporting or display lists of resource trees and forks. The CMDB has to be able to visibly display its data in ways that let humans use the information in impact assessments for Change Management, privilege determination at the Service Desk, troubleshooting by Incident or Problem Management, and dozens of other ad hoc inquiries from all over IT.

This requirement goes way the simple "directory tree" listings so commonly found in most alleged CMDB products today. Federation, reconciliation, and synchronization is worthless if a user in IT cannot get a definitive, understandable answer to their complex questions quickly. This requires representing complex relationships between CIs graphically, on demand.

Summary

An enterprise CMDB is not a database, it is a complex software system that has to federate other data stores, reconcile alternate views of the same data, detect unauthorized changes, synchronize approved changes with its own metadata store, and be able to dynamically represent configurations graphically on demand. This is no small task, and also the reason there are so few true CMDB solutions available today.

As you go forward with your own CMDB plans, keep the concepts of federation, reconciliation, synchronization, and modeling in clear view. Dig into them to understand them. If you are in the throws of purchasing a CMDB product, ask your vendors how they handle these issues. If you are building your own CMDB solution, ask your developers how they plan to accomplish these tasks. In all cases make sure you create processes to monitor and ensure that federation, reconciliation, synchronization, and modeling occur. Failure to manage these critical issues can quickly convert your CMDB project from an asset to a liability.

Forewarned is forearmed! Now you can at least "talk the talk as you walk the walk!"



[Subscribe](#) Its Free!



[RSS](#) Stay up to date!



[PDF](#) Pass it around!

Where to go from here:

- [digg](#) (discuss or comment) on this article. Show your support for DITY!
- Subscribe to our newsletter and get new skills delivered right to your Inbox, [click here](#).
- Download this article in PDF format for use at your own convenience, [click here](#).
- Use your favorite RSS reader to stay up to date, [click here](#).

Related articles:

- [Configuration Management for the Rest of Us](#) for more on establishing Configuration Management.
- Browse back-issues of the DITY Newsletter, [click here](#).

