

Future CMDB

By [Hank Marquis](#)

Hank is EVP of Knowledge Management at Universal Solutions Group, and Founder and Director of NABSM.ORG. Contact Hank by email at hank.marquis@usgct.com. View Hank's blog at www.hankmarquis.info.



ITSM Industry powerhouses have stopped fighting and are now working together. They recently released a joint whitepaper describing how they plan to work together to solve the thorniest CMDB technological hurdles...

Anyone who has read this column over time knows that I think a CMDB is more of a knowledge management system than an IT Asset Manager (ITAM). In previous issues of this newsletter I described how a CMDB solution, as a knowledge management system, has to federate data from many sources.

Federation is basically collecting data from many heterogeneous sources including many systems and network management platforms, log files, server files, router data, etc. The issue with federation is simple: every vendor has their own interface -- if they have an interface at all.

While CTO of Opticom, an ITSM software provider, our greatest technical hurdle was connecting to the diverse range of systems, collecting data stored in many formats, and then rationalizing that data to our own internal format.

My experience is not unique. Any vendor or practitioner about to embark on a CMDB creation project quickly runs smack into the brick wall that is federation.

But the future holds promise. A white paper release this week from BMC, CA, Fujitsu, HP, IBM, and Microsoft details how these powerhouse vendors plan to provide a common interface for CMDB products specifically for the purpose of federation.

Following I give my analysis of what I think this announcement really means to us average IT people trying to "do ITIL".

The Issue

As the whitepaper describes, there are many Management Data Repositories (MDRs) in an enterprise today: system management tools, service level management tools, application discovery systems, mainframes, server management systems, and hardware management systems. Each vendor has its own MDR, and every MDR is in its own format.

This is what makes data federation so difficult. It also has a dramatic impact on reconciliation. Reconciliation is understanding the two views of the same CI from two different management systems is really the same CI. The two issues go hand in glove. The proposed architecture addresses both. Synchronization is the ability to detect changes in federated data and generate an alert, and modeling is the ability to query the CMDB as required. The white paper addresses all of these critical issues.

Interestingly, the example in the whitepaper describes the utility of an "open" approach to CMDB federation (and reconciliation) from the viewpoint of the service desk. Part of the whitepaper presents a data model overview which addresses the prickly issue of reconciliation. This is very good news because it should address naming conventions and resource identification methods.

The whitepaper proposes four (4) CMDB Services.

1. Administration of the Federated CMDB (Federation)
2. Resource Federation and Registration (Reconciliation)
3. Resource Query (Modeling)

4. Subscription and Notification (Synchronization)

Administration of the Federated CMDB -- Federation

The proposed *administration service* provides the primary interface for a Federated CMDB. All MDRs participating in a Federated CMDB register with the administration service, and each MDR associates with schemas and capabilities that it supports. The administration services allow management of the federated data.

Resource Federation and Registration -- Reconciliation

A proposed *federation service* provides the interface to add, modify, and delete resource definitions. The federated CMDB manages a directory of all registered resources (configuration items, relationships, and process artifacts like incident and change records). It is responsible for reconciling the different perspectives of a resource's identity into a common "resource context". It may also manage additional data properties provided by MDRs when resources are added to the federation service.

Resource Query -- Modeling (Mapping & Visualization)

A proposed *resource query interface* allows client applications to retrieve data from the Federated CMDB based on certain criteria. This includes an instance query to retrieve items based on identifying properties, and a relationship query to enumerate a set of items and relationships based on attributes of the relationships and endpoints.

Subscription and Notification -- Synchronization

A proposed *subscription service* provides applications the ability to register for data change events that could occur in the Federated CMDB. Applications needing to respond to data change (update, insert, delete) events will subscribe to that event and will be delivered notification when such an event occurs.

Analysis: the Good, the Bad, and the Ugly

The very fact that vendors who compete on a daily basis are working together shows that it is very, very difficult to achieve federation. The good news is that they are working to establish a so-called "industry standard" by which vendors can interoperate. This is very, very good. This can make our lives in IT much easier. If you are a vendor, it makes your products more useful. The more sources of data in the CMDB the better the decisions we can make in IT. With a standard interface we won't have to spend so much time and money hand-crafting interfaces and polishing data structures. It is good that the paper addresses the four key elements I discuss in ["Enterprise CMDB", DITY vol. 6, issue 17](#).

The bad news, which does not have to come to pass, is around the potential issues with this approach -- simply that it is NOT a standard. We have all seen how vendors, constrained by "standards" make subtle implementation modifications. If not managed very carefully, this results in a tower of babble. Since this is not an ISO (or any other standard) there is absolutely no requirement for any vendor to follow it, or to follow it rigorously. Additionally, where is Cisco in all of this? With 86% of the market for networking hardware last time I looked, why are they absent? And what about the other vendors? Where are the midrange vendors like TouchPaper, InfoVista, TeamQuest, and FrontRange? An "open standard" pushed by a select group of companies is bound to engender resentment on those "outside" and maybe even alternate solutions. For a truly open standard, this ought to be pushed off into ISO or some other standard making body for ratification and control.

The ugly issue not addressed anywhere is Microsoft's attempts to patent the CMDB. As chronicled by the IT Skeptic, Microsoft patent application number 20060004875, in which Microsoft has applied for a patent on "CMDB Schema" that "includes an entity to store information identifying configuration items and a separate entity to store attributes of the configuration items. The CMDB schema may also include a separate entity to track relationships between configuration items." Sounds like Microsoft at least understands where the CMDB is heading. You can view the patent application [here](#).

This is a direct attempt by Microsoft to patent most of the techniques expressed in this whitepaper. I would love to have been a fly on the wall to hear how they rationalized this patent application to the other CMDBF members.

We have to take a wait and see attitude on this one -- see what the CMDBF actually produces. You can read the whitepaper here: [CMDB-Federation-white-paper-vision-v1 0.pdf](#).

