

# **THE FEDERATED CMDB**

## **Three Applications of the Term**

**Version** : V 1.0  
**Date** : November 10 2006  
**Author** : Troy DuMoulin



## 1 EXECUTIVE SUMMARY

The emerging capability of federating data sources is the biggest boon and the largest potential pitfall that has arisen for the discipline and process ITIL calls Configuration Management!

On the surface this may appear to be a contradictory statement. However, as is the case with all good intentions, it all comes down to the application. First and foremost, database federation is an absolute must for a successful Configuration Management Database CMDB implementation.



## Table Of Contents

1	EXECUTIVE SUMMARY.....	2
2	THREE APPLICATIONS OF CMDB FEDERATION .....	4
2.1	Application 1: CMDB Federation Is The Exception Not The Rule.....	5
2.2	Application 2: CMDB Federation Is The Rule Not The Exception.....	8
2.3	Application #3: The Dynamic View From Existing Distributed Data.....	10
3	SUMMARY .....	11
4	ABOUT PINK ELEPHANT.....	12



## 2 THREE APPLICATIONS OF CMDB FEDERATION

One, and perhaps the most fundamental definition of a federated CMDB is that of a master database (aka the CMDB) that references an external database for a subset of information. Simply put, when I am looking at a record representing a Configuration Item (CI) some of the information I find in the record is stored in the central CMDB and some of the information is stored in an external source or database. The key is to make this appear completely transparent to the person accessing and reviewing the central CMDB record

While federation is important, the challenge we face is that there are three different uses or applications of the term. Of these three uses, I personally believe two of them represent strategies for data consolidation that are likely roads to long term failure.

The three applications of the term ‘federation’ are:

### 1) **CMDB Federation is the exception, not the rule.**

The CMDB is a tool used by all IT Management functions and processes and the effort is made to *consolidate* and centralize most, if not all, the “like” data sources into the CMDB. We will link to other data sources only where it makes sense

### 2) **CMDB Federation is the rule, not the exception.**

The CMDB is a tool used only for service management and there is no reason for data source consolidation. We will leave everything where it is, cherry pick and replicate the attributes and data we want into a centralized CMDB record and make sure they are up to date through reconciliation.

### 3) **CMDB Federation is the dynamic, spontaneous creation of a CI view from different existing data sources.**

The “CMDB data” already exists in the myriad of data sources and there is no need to create a central record – we just need to pull it all together in a virtual data warehouse model to create a current view of the existing data sources. In this approach, a central record is never created.

Since I am obviously biased, I will start with the application of the term ‘federation’ that I believe to be the most appropriate use of the term, and then carry discussion on through to the least appropriate.



## 2.1 Application 1: CMDB Federation Is The Exception, Not The Rule

This approach starts with a belief in four core truths:

- 1) The CMDB is the central repository of record for all IT Assets.
- 2) The CMDB is the singular trusted source of data and is used by all IT Management disciplines and not just ITIL processes.
- 3) Federation to other databases applies only when the child data is residing in an IT tool or application database that has unique functionality and the source can be trusted as accurate and under control.
- 4) A best practice of data management is that you only store and manage it once if at all possible.

The outcome of these four truths or principles implies that there is an overall goal to *consolidate* all of the IT data management sources (regardless of who manages them) into one master source when those sources are 'like', meaning that they have the same basic purpose (i.e. store data about CIs).

**Additional Consideration:** The number of attributes stored in a source database being evaluated for consolidation does not qualify as a basis for unique functionality. A decision on how many attributes should be managed in the CMDB should represent the balanced interest of all stakeholders interested in the data. This includes the owners of the original source database as well as the other groups and processes that require access to it. The final decision about what attributes will be managed at a controlled or un-controlled level is a deliverable of the CMDB data modeling step of the project.

All of these like-data sources, be they spreadsheets or enterprise database solutions, are candidates for consolidation into one source, the central CMDB. And yes, by consolidation, I mean elimination of the original database in favor of the central one. Remember the principle of managing data only once.

However, there are many sources of data that have unique functionality. For example; data stored in Active Directory which is used for rights management, a systems management tool used for discovery or event correlation, or a financial management application used for costing and billing.

Good Examples of Federation include:

1. When I look at a CI record and examine the financial attributes I am actually looking at data stored in the Enterprise Resource Planning application.



2. When I look at the hardware attributes, these are sourced from a discovery tool where we have selected only 12 attributes for the CI record out of a possible number of hundreds that are discoverable from the existing source.
3. People and organizational data is coming from Active Directory, the HR tool or the email global address list.

Each of these examples represent a source of data with a unique functionality which is above and beyond the goal of managing data about CIs and represents an excellent candidate for federation rather than consolidation.

While it is generally considered a best practice to manage data once, it is not always possible to avoid data replication. Often for technical or security reasons it is not possible or desirable to provide a real time view into the child data source; therefore, it is essential to have the capability to perform a regular reconciliation of the two sources through an automated means.

### **2.1.1 Three Traditional Objections To This View Of Federation**

- 1) The ITSM processes don't need all of that data, so why would I mess up the CMDB with information these processes don't require?
- 2) The database would be huge and performance would be an issue.
- 3) The cost of this consolidation would outweigh the benefit.

The first argument assumes that having a source of record where we track all relevant CIs in relationship to each other and IT services is only of use to the processes focused on service management. This is simply not true – procurement, software configuration, audit, security architecture and engineering, and project management to name a few also have a use for the CMDB.

The second argument is not well-founded since it is unclear how many CIs we are actually talking about managing. Many organizations by the definition of an asset deem certain CIs not worthy of management simply due to their relative financial cost (ie: keyboards or perhaps monitors). Larger organizations can get into the double and maybe triple digit thousands, but that would be the outside number and this is well within the capability of the enterprise database solutions in use today.

The primary reason why an organization has multiple solutions for managing data is a result of history, politics and IT procurement practices focused at the domains. Based on a traditional technology management view, each IT domain is managed as a unique function and procures tools for its own needs (e.g.: the database group has a database on databases, the Server group has one database for Unix boxes and another for Wintel Machines, the application groups track their applications, the network group is just



concerned with network components, etc.) From this perspective, each group has built separate data sources to manage their own CIs.

Unless an IT organization can understand how any technology component connects to another and how they both impact a business process, it is very difficult to claim alignment with business objectives. What do you do when you realize that managing each domain in mythical isolation prohibits you from understanding the relationship of dependency between them? It is only when an organization begins to move to a service orientation that this question becomes a burning issue. This paradigm shift requires the creation of a CMDB where CIs can be modeled in relationship to each other. Even more important is how the CIs are bundled to create systems that in turn support IT Services as consumed by the business.

Once upon a time the business had separate applications to support business processes such accounts receivable, inventory management, procurement, payroll, etc... Each of them had their own separate databases on different platforms which needed to be connected through complex integrations. However, IT stepped up to the plate and said wouldn't it make sense to have a suite of connected applications all getting their data from the same primary central source with federation to other internal and external sources applied when necessary.

*The Enterprise Resource Planning (ERP) Suite was born!*

If we preach this as a good strategy for the business, what makes IT any different? The answer of course is that we are not – instead, we are only a few years earlier in the management maturity curve.

The follow-on question then becomes why keep two tools to maintain data about the same CIs? Unless there is some significant functionality difference the choice to not consolidate these duplicate sources of CI data at some point in the near future is based on politics and emotion not logic and cost.

The final argument assumes that managing data once in a single repository is somehow more expensive than managing and supporting multiple redundant data sources and tools. It would be a very interesting exercise to do a Total Cost of Ownership assessment on the various tools in the organization today and compare it to a centralized solution. For many organizations this model would represent a significant consolidation savings through the removal of redundant tools, streamlining the associated management processes and a reduction in the resources required to support them.

Each time an organization buys hardware, installs, supports and buys licenses and maintenance contracts for data management solutions it incurs an initial and ongoing cost. To consolidate even a third of these could represent a significant long term savings.



To be fair for some organizations many of the current data sources are low tech and have limited overhead costs. It is difficult to suggest that doing away with the spreadsheets in favor of an enterprise CMDB would save money based on consolidation: however it is a true statement that maintaining both the spreadsheet and the central CMDB is not an efficient use of resources. Add to this the cost of developing complex integrations that are not required and the third argument is hard to justify.

## 2.2 Application 2: CMDB Federation Is The Rule, Not The Exception

The second application of the term Federation is the reverse of the previous one. This use assumes a premise that the CMDB is used only by ITIL processes. Thus, it also largely presumes that there is limited value in trying to remove redundant data management tools in the IT organization.

In this model, the siloed IT functions graciously permit “their” respective data sources to be tapped into by the central CMDB for which they feel limited to no personal accountability. In this model the CMDB is tolerated as required to support management processes like Change and Incident instead of being considered the central repository of record for general IT Management.

In concept what is being suggested in this model is the creation of an instantiated data warehouse where source data is drawn to a central record that gathers history over time.

*A data warehouse is a database geared towards the business intelligence requirements of an organization. The data warehouse integrates data from the various operational systems and is typically loaded from these systems at regular intervals. Data warehouses contain historical information that enables analysis of business performance over time.*

[www.oranz.co.uk/glossary\\_text.htm](http://www.oranz.co.uk/glossary_text.htm)

There is a primary cultural difference between ITIL Configuration Management and distributed Inventory Management. Distributed Inventory Management represents a principle that each group updates their data in accordance with their personal need and culture. For some groups like the Unix Admins, this might be a daily activity, but for others it could be weekly, quarterly or based on best effort.

ITIL Configuration Management on the other hand requires almost real time accuracy in order to be trusted to the point people will use this data for various processes and management decision making.

Given that, the challenge with this second approach to federation is that it is very difficult to enforce data update and management at the source. If there are challenges with the controls managing the original source database, the owners of an ITIL focused CMDB



## The Federated CMDB

---

are out of luck and they have little to say about it (you get what you get, and be thankful you get it!)

In addition the management overhead with this approach is much higher than the previous model. Every CI added to the environment would always needs to be added to both the source database and the central CMDB with the additional effort of establishing the link and reconciliation rules between the two sources.

Please consider that many organizations have a high volume of new CIs added to the environment on a regular basis. The creation of a backlog and the possibility of failure based on relational and technical complexities are exponentially higher with this model. Several organizations I have spoken to over the years have begun with this approach since it was politically easier of the two models, but they have had to re-start their project (with the associated cost implications) using Application #1: Federation as the exception, not the rule.



### 2.3 Application #3: The Dynamic View From Existing Distributed Data

The third definition of federation is much like the one we just addressed and struggles with many of the same challenges; however, in this application a central CI record is not actually created. Instead, this model is based on the concept of creating a dynamic on-demand compilation of a view from existing distributed data each time a request is made.

This model is a sometimes referred to as a Virtual Data Warehouse

*An enterprise data warehouse constructed of multiple data marts and a request broker computer application. The data warehouse does not physically exist except through out the formation of the integrated data marts. [University of Southern California](#)*

While this solves the problem of creating two sources for managing the same data, it comes with another inherent challenge. In the ITSM community, it is considered best practice to not only create relationships between CI records but also relationships from CI records to process records.

For example, the following relationships are critical to support ITSM:

- The ability to associate CI records to each other to model systems and services as consumed by the business.
- The ability to associate Incident, Problem, Change and Release records to CIs throughout their lifecycle
- The ability to associate specific SLAs and business rules to CI records

When the CI is not instantiated and is only a virtual view on distributed data sources, these CI-to-Process relationships become severely restricted if not impossible.



### 3 SUMMARY

A key benefit of adopting a best practice standard is the adoption of a common language. This is without a doubt one of the things many organizations find attractive about ITIL; however, consider that the term 'federation' is used by many in regards to implementing the CMDB but has three separate applications depending on whom you are talking with.

The next time this topic comes up, it would be wise to conduct a definition check as to which one of the three applications is meant.

To help you assess the right approach for your organization, you should:

- 1) Review all current repositories and data sources.
- 2) Determine if these sources are "unique".
- 3) Define the long-term CMDB scope to include all data sources that are not considered unique or are not in a controlled state (this can be phased over time).
- 4) Identify owners of these data sources and include them in the overall CMDB planning and design initiative.
- 5) Document and socialize the planned federation approach agreed to by the design team.
- 6) Develop an Object Model that incorporates this approach and the required attributes and relationships that will add value to your CMDB.

Creating a federated data model is a necessity for the successful implementation of a CMDB strategy; however, the approach you use can make a difference between long term success and failure. By experience, we find the first approach listed in this paper is by far the most successful over the long term even though it is harder to politically initiate and achieve without a clear vision and definitive executive support.

**Data source consolidation and centralization is the rule and federation is applied only when a business case for unique functionality has been established!**



## 4 ABOUT PINK ELEPHANT

Pink Elephant ([www.pinkelephant.com](http://www.pinkelephant.com)) is the world leader in IT management best practices, offering conference, education and consulting services to public and private businesses globally, and many listed in the Fortune 500. The company specializes in improving the quality of IT services through the application of recognized best practice frameworks, including the Information Technology Infrastructure Library (ITIL®).

### Service Lines

Pink Elephant's service lines each provide different, but complementary business solutions:

- *Business Process Consulting*: Using the ITIL best practices approach to IT service management as a springboard, Pink Elephant provides end-to-end solutions – from assessments, to strategic planning to implementation, continuous improvement and beyond. Experienced consultants work hand-in-hand with customers every step of the way
- *Conferences & Special Events*: Pink Elephant is the world's largest producer of IT service management conferences and delivers several major events per year
- *Education*: Pink Elephant is the most prolific creator and widespread distributor of ITIL training, delivering three levels of certification – Foundation, Practitioner and Management. Pink Elephant is internationally accredited by independent examination institutes that manage the ITIL certification program worldwide
- *ATLAS™*: ATLAS™ is a secure, web-enabled knowledge management system containing much of Pink's highly valued intellectual property – ready and waiting for users to access, copy, customize and re-use

### ITIL Leadership

Pink Elephant has grown to become recognized globally as *The ITIL Experts* and is very proud of its commitment to the ITIL best practice framework. In fact, Pink Elephant has been involved in the "ITIL project" since its inception in 1987, and was recently selected by the UK's Office of Government Commerce (OGC) to help author ITIL's next version. Furthermore, Pink Elephant:

- Supported the development of ITIL's core books:
  - Service Support (English and French editions)
  - Service Delivery (English and French editions)
- Launched PinkVerify™, the only independent certification program worldwide that recognizes software that supports specific IT management processes
- Created the International IT Service Management Conference & Exhibition, one of the largest events worldwide solely dedicated to ITIL



- Facilitated plans for an ITIL examination centre in North America (Loyalist College in Belleville, Ontario, Canada)
- Introduced ITIL internationally to companies across a wide variety of industries, sizes, technical platforms and corporate cultures
- First offered the Foundation, Practitioner and Management ITIL certification levels publicly in North America
- Was a founding member of the IT Service Management Forum (now *itSMF*) – the worldwide networking group for IT service management professionals

### **Awards**

Pink Elephant is recognized as a progressive and successful company and is the recipient of the following awards that reflect its corporate leadership excellence and business results:

- Top 100 Woman Entrepreneurs – 2001 – 2006: Awarded to Pink Elephant CEO, Fatima Cabral, by PROFIT magazine
- Ontario Global Traders Award – 2005: Awarded by the Ontario Government for achievements in innovation, leadership, product excellence and expansion into new markets
- Business Excellence Award – Awarded by the Federation of Portuguese-Canadian Business and Professionals to Pink Elephant CEO, Fatima Cabral
- Top 100 Fastest Growing Companies in Canada – Awarded annually by PROFIT Magazine. Pink Elephant was recognized as one of Canada's fastest-growing companies (based on a comparison of revenue growth for five consecutive years)
- Top 100 Canadian IT Professional Services Organizations – Awarded by Branham300
- EXIN Award: Given to the organization with outstanding achievements in promoting the IT Service Management framework outlined in the IT Infrastructure Library (ITIL) – the world's most popular set of IT management best practices

To learn more about Pink Elephant's services visit [www.pinkelephant.com](http://www.pinkelephant.com), or call us at 1-888-273-PINK.

*ITIL® is a Registered Trade Mark and a Registered Community Trade Mark of the Office of Government Commerce, and is Registered in the US Patent and Trademark Office.*