



BMC ATRIUM CMDB – DATA POPULATION

INTRODUCTION

The Configuration Management Database is the foundation of an Information Technology Service Management (ITSM) enterprise solution. Many organizations do not have a formal or rigorously maintained Configuration Management Database (CMDB), or their CMDB data is incomplete, inaccurate, or out of date. A CMDB that contains incomplete or inaccurate configuration items (CIs) and asset data can render the IT Service Management function ineffective. The absence of a centralized definitive source of reliable and accurate CI and asset data can result in many problems including for example:

- Customers, users, and leaders question the validity of CMDB information needed to make business decisions regarding the IT infrastructure and the services provided
- Determining the full impact and costs associated with a proposed change to the infrastructure becomes virtually impossible
- Costs associated with software license renewal, hardware refresh, and maintenance fees are difficult to forecast resulting in inaccurate budgets and cost overruns
- Considerable time is spent on physical audits of asset inventory and verifying or re-verifying CMDB data that may be out of date the next day
- Determining the impact and cost an outage has on a specific service is more difficult to determine

G2SF's ITSM subject matter experts have perfected a highly effective method of populating and accurately maintaining the BMC Atrium CMDB. This method has many benefits including:

- Establishing a central repository of definitive data to accurately control and manage assets throughout their lifecycle
- Minimizing the risks and potential downtime associated with introducing and releasing changes into the production environment
- Providing reliable data regarding all assets (CIs), their relationship to key business services, as well as their impact in the event of a failure or incident
- Maintaining an accurate and up to date inventory of all assets and their locations to facilitate Tech Refresh
- Providing effective Software License Management to ensure costs are properly managed and budgeted
- Leveraging the use of automated workflows and repeatable processes to ensure that the CMDB contains the most accurate and reliable information available

This white paper provides a high-level overview of the key steps that were taken to successfully develop, populate, and accurately maintain a reliable enterprise CMDB within a large and complex DoD global organization. The key steps include:

1. Building the System Foundation
2. Integrating with External Data Sources
3. Normalizing and Reconciling Data

This white paper is intended to be used by G2SF exclusively for informational purposes only and shall not be copied or reproduced in whole or part without G2SF's written consent from BD@G2SF.com.



BUILDING THE SYSTEM FOUNDATION

Establishing key data elements from a business perspective is critical. The definitive source of business data should be a business system that is intentionally used for accurately collecting and maintaining source data in one single repository, such as the Atrium CMDB. Rigorous processes should be established to ensure data integrity. The keys to successfully populating a CMDB revolve around identifying and addressing critical elements such as building a solid system foundation including defining an enterprise Common Data Model (CDM), designing integration with external discovery sources, and developing and implementing normalization and reconciliation rules.

Accurate and Complete Helix (Remedy) Foundation Data

G2SF's initial CMDB development efforts focuses on identifying and building Atrium foundation data to align with the customer's organizational structure. A properly designed ITSM platform allows the business to efficiently manage IT services, assets, costs, and risks to make effective decisions regarding the current and future use of technology to meet mission goals and objectives. The goal in properly designing the CMDB with accurate foundation data is to ensure a common understanding of how the CMDB and the ITSM platform will be used to meet ITSM and mission goals. The ITSM platform is synchronized with key data elements used by customers. Synchronizing the key data across the enterprise not only allows IT to provide services in a way that makes sense to the mission, but also allows IT to report back to leadership in a context they understand. The resulting foundational data model represents an alignment between the organizational structure and the IT support structure. The Atrium foundation data is the cornerstone of all system workflows including:

- Auto ticket routing. This workflow automatically moves/routes a ticket to the next assignment group in the process once all essential criteria have been captured.
- Approval rules. This workflow checks pre-defined conditions that must be met to approve a change or request. Only specified individuals have authority to approve based on ticket type and criteria.
- Reporting and metrics. This workflow shows who is being supported (Organization) and who is providing the support (IT Support Structure).
- Entitlements. This workflow manages who can view services offered on the service catalog and who can request those services.
- Financial tracking. This workflow allows for tracking financial data associated with a specific organization, company, or department.
- Asset management down to the organization structure. This workflow allows for tracking assets associated with a specific organization.
- UIC Admin Console. This workflow allows Asset Admins to manage relationships and workflows specific to their UIC.

While identifying and restructuring the foundation data, G2SF recognizes the importance of the Unit Identification Code (UIC) used by DoD organizations, and as a result has developed a custom console that allows users to manage foundation data without having to access the foundation data tables. The UIC is a key foundation element used by the ITSM system for product ordering, asset tracking, and CI recording within the CMDB. The UIC data is used to maintain - managed by, supported by, owned by, and approved by company, organizations, and department attribute mappings.



The Unit Identification Code (UIC) Administration Console is a custom-developed console and workflow that enables the users to manage foundation data UICs used to create relationships between the owning and supporting organizations for each UIC. The workflow automates the population of CIs in the CMDB with pre-mapped owning organizations and support group relationships during the normalization process. This workflow automation provided a G2SF client with a 32% reduction in man-hours spent manually populating user managed information within the CMDB. The UIC Admin Console provides DoD support staff the ability to comprehensively report on hardware asset inventory to 1) quickly determine the owner of an asset when performing refresh and auditing activities, 2) provide Asset Managers the ability to search for CIs by support group and role, 3) allows Asset Managers to add the "View Asset List by my Support Group" widget to their Remedy Homepage to display the CIs that belong to their support groups by role allowing quick access to the CIs they are tasked with supporting, 4) provides support group relationships allowing support staff quick access to view the support group members, people records, and information on their support roles in relation to the CI directly from a CI record.

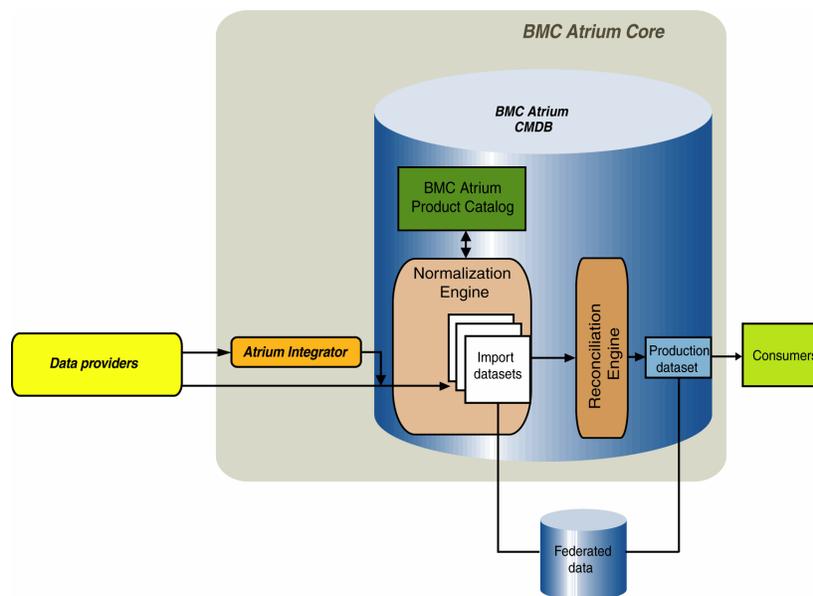
Common Data Model (CDM)

Once all foundation data is identified and structured, G2SF focuses on modeling the CMDB. G2SF conducts workshops on data modeling, mapping, scope, and requirements determination that result in the design, development and implementation of the Atrium CMDB used to store CI data and CDM information in a manner that allows users to build and view the relationships between CIs. To enhance the relational service model, G2SF engineers also create a relationship schema that links CIs to specific services. This allows decision makers to accurately assess user and IT impacts of incidents and proposed changes. Every Configuration Management effort must begin with data modeling. G2SF works with our customers to create an Enterprise Common Data Model that contains CI definitions and provides appropriate policy and direction regarding the types of CIs to be included or excluded from the CMDB. The objective is to maintain a balance between providing relevant CIs and CI information necessary to ensure maximum service availability without including a level of detail that is of limited value and that may create CI maintenance challenges. G2SF ensures that each configuration item type is properly identified and represented with the base class and attributes necessary to perform Incident Management, Change Management, Configuration Management, Problem Management, Asset Management, and Knowledge Management activities. From the information obtained, a list of CI types is created identifying common attributes for all identified CI types and unique attributes for each specific CI type. Using customer approved attributes, G2SF uses templates to define discovery information needed to populate the CMDB.

Once the CDM is documented and approved, G2SF assesses current asset databases in use and designs the process to migrate identified data and attributes into the Atrium CMDB. This data migration consists of automated and/or manual exports of data into predefined spreadsheets and performing data field mapping to complete the data migration. Once all the data is successfully migrated, G2SF sets the baseline of the Production Data Set or the Atrium CMDB.

INTEGRATING WITH EXTERNAL DATA SOURCES

Integrations with auto discovery tools like Bell Arch BellManage, Microsoft SCCM, HP Universal Discovery, and BMC Discovery were developed to populate auto discovered CI data. G2SF utilizes both out of box adaptors and custom developed adaptors to establish connectivity with external data sources. Integration adaptors are used to retrieve data from external discovery tools, transform source data into data that can be used by the Atrium CMDB, and populate source data into defined datasets in the Atrium CMDB. Data discovered by external discovery sources is then ready to be normalized, reconciled, and promoted to the Production Dataset and is then available for use by other BMC modules such as Incident Management, Asset Management, and Change Management that share the BMC Atrium data. The BMC Atrium Integration Engine (AIE) product enables G2SF engineers to transfer data between an external data store and the BMC Atrium Configuration Management function. The diagram below depicts the data flow.



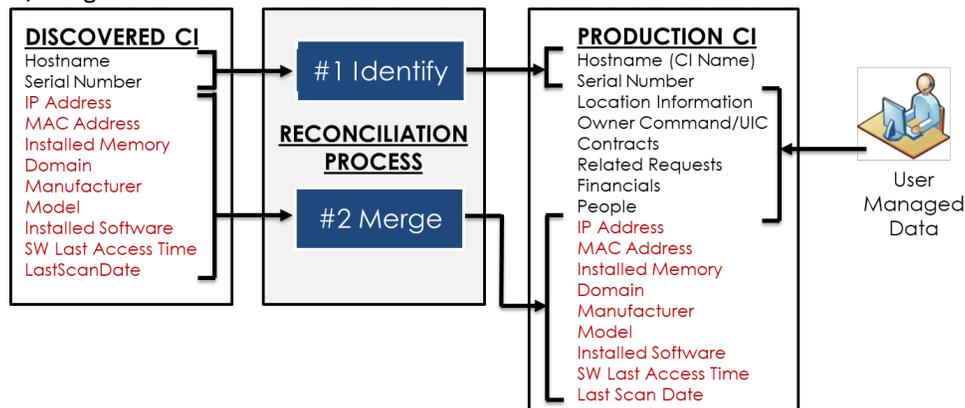
The basic steps necessary to populate the CMDB with data from external discovery sources are as follows:

1. Review the data source's database table and view structure and build the desired database queries. This step entails conducting a full review of the data source's database to validate the data stored in the database and to ensure the proper tables/views are being queried. This is necessary because in most databases, data for one configuration item is stored across multiple tables/views. Ensuring the correct table/view is selected is paramount to making sure the data accurately represents the configuration items necessary to populate the CMDB. Once the proper tables/views are determined, the database queries need to be built to retrieve and import data into the CMDB.
2. Build scripts and transformations (using Pentaho Spoon) that will retrieve the data and populate the Atrium CMDB. This requires all data to be carefully analyzed for discrepancies prior to import.
3. Load the external data into staging dataset in the CMDB and analyze the configuration items to determine if the product catalog needs to be updated. Product Catalog entries will need to be created and updated to ensure the new configuration items coming from the external source can be normalized. Several normalization aliases will need to be built to ensure all similar type configuration items have the same nomenclature. This is essential for CMDB reporting once the data is loaded into the gold dataset within the CMDB.

NORMALIZING AND RECONCILING DATA

G2SF develops normalization and reconciliation rules to ensure the accuracy of the production data set within the Atrium CMDB. Normalization is used to merge discovered data with non-discovered data and to promote the data to the production data set. Reconciliation is used to flag data changes and ensure no duplication of data.

The graphic below illustrates in greater detail which Discovered CI data attributes are merged with the associated Production CI data record. The attributes identified during foundational and common data model workshops are researched to determine where the data enters the system. The Reconciliation Process identifies the Discovered CI relates (ensuring accuracy) to the Production CI by matching specific attributes, then appends/merges the Production CI record with the additional attributes of the Discovered CI.



The following steps are performed in the Normalization and Reconciliation process:

1. Normalize the new data in the CMDB using the normalization engine that will modify configuration items and their relationships to standardized records. The normalization engine updates the Manufacturer, Product Name and categorizations on the CIs to ensure they are consistent and that they match what is in the Remedy Product Catalog.
2. Reconcile the data in the Remedy CMDB using the Reconciliation Engine. This is a two-step process that includes an identification stage and then a merge.
 - a. During the identification stage, the system performs a lookup in the gold dataset with the data in the staging dataset using configurable identification rules. These rules are in place so the reconciliation knows what attributes on the CI are used to uniquely identify one CI from another. When the Reconciliation Engine finds a match in the gold dataset, it then links the record to the one in the staging dataset. Once the CIs are linked, any changes on the staging record can be merged to the gold dataset CI.
 - b. Once the data has been identified, the next step is to merge the data. Based on configurable rules in the CMDB, data from the staging dataset is promoted to the gold dataset. Merge precedence rules are setup so that only specific attributes selected will have their data promoted. Otherwise, more weight is given to the attributes on the record in the gold dataset. After the data is merged and updated, users are then able to view and report on accurate information regarding a given CI.

Together, the CMDB, CDM, Atrium foundation data, and normalization and reconciliation rules are used to accurately track assets throughout their lifecycle and to determine how the assets are being used in support of the business or mission. It is the process of identifying, collecting, mapping the Discovered CI to append the Production CI record in the CMDB that provides value to the customer. This automated process ensures all customer-identified CI information is reliable, accurate, located in a single location, and easily accessible to authorized users to make better and faster decisions.



ITSM SERVICES AND SOLUTIONS

As the foundation of an Information Technology Service Management (ITSM) enterprise solution, there are several tangible benefits customers will realize from G2SF's involvement in designing, developing, populating and managing an accurate and reliable Atrium CMDB. Other ITSM related services G2SF offers includes for example:

- Baseline ITSM Maturity Assessment: identification of organizational inefficiencies and opportunities for measurable improvements that will reduce costs and increase customer satisfaction
- ITSM Software Selection, Implementation, Integration, and Management: fully integrated tool suite that automates processes and business outcomes saving time and money
- IT Business Process Definition, Design, Implementation and Reengineering: repeatable processes that are routinely used due to efficiency, simplicity, and time savings
- Migration of ITSM to the cloud to facilitate ITSM as a managed service
- IT Service (Help) Desk Consolidation, Management, and Support: centralized single point of contact for all service requests and incident resolution resulting in a reduction in operating costs, improvements in service, and organizational effectiveness
- Implementation of Artificial Intelligence (AI) and Robotic Process Automation (RPA) technology to facilitate monitoring, routing, tracking and resolution of tickets and incidents, including security related threats and incidents
- Integrating use of the CMDB with the Cybersecurity function to provide real time information on vulnerable assets
- Organizational Change Management: rapid adoption of changes that positively impact the organization and its mission
- ITSM Education and Certification Training: common understanding of standard operating procedures, processes and tools

G2SF POINT OF CONTACT

For additional information on this topic please contact:

Dan Nichols, Sr. Program Manager and IT Service Management Practice Lead at bd@g2sf.com.