ACTIVATING CONFIGURATION MANAGEMENT IN AN IT ORGANIZATION

White Paper

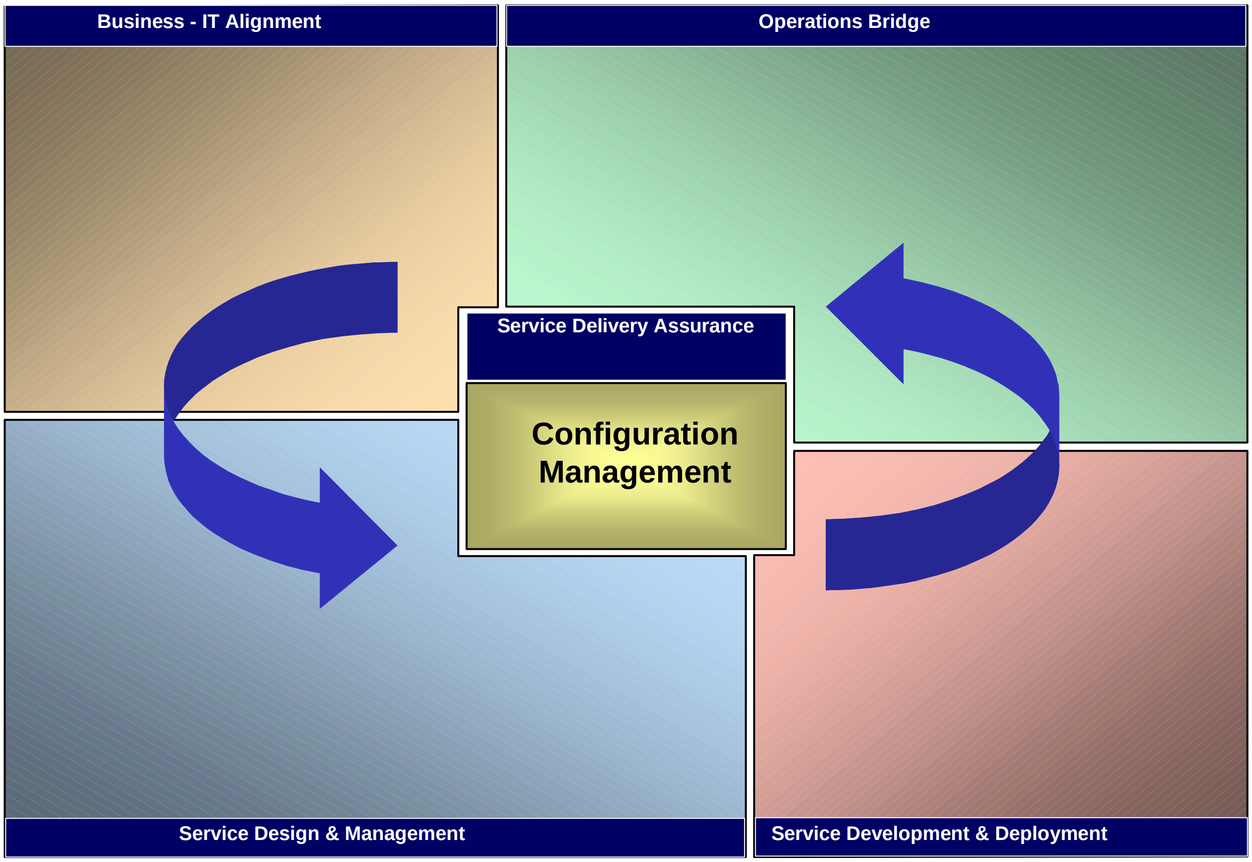

People + Process + Technology

invent

Acknowledgements

The development of this white paper was a team effort, involving many people from different HP entities and countries. I'd like to take this moment to identify each team member by name in order to recognize the value of his or her contributions:

Special thanks to Vladimir Kufner, David Mariscal, Kathleen Wilson-Edwards, Sev Canzona, Derek Gillard, Peter Lijnse, Jeroen Bronkhorst, Ruud Ligtenberg, Jeroen Wiebolt, Paul Miller, Philip Sousa, Lindsay Parker, Armin Hampel, Milan Hojdar, Robert Krajca, Gerry Machado, Gavin Runnalls, Klaas Zweepe and Gerben Verstraete. Thanks also to Tjerk Feenstra for his original idea for this white paper and to Regina Kershner and Rick Owen for their support in the development of this paper.

Abbey L. Wiltse
IT Management Process Consultant
Hewlett-Packard Consulting (Canada)

TABLE OF CONTENTS

1. CONFIGURATION MANAGEMENT PROCESS ESSENTIALS

1.1 OVERVIEW

It is fundamental that IT organizations have the ability to track, record and report on IT infrastructure data. Organizations have attempted over the years to get a handle on this aspect of their business. In the early years of IT, the process to keep this information up-to-date and accessible to all was simple. You were most often dealing with a single technology platform that provided it's own configuration files and this information was accessible by all that needed to conduct business using this data.

In today's complex world, IT organizations experience the reality of managing multiple technology platforms across many technical domains, of performing in a permanent state of change, of responding to increased business demands and of reacting to the fast-pace of mergers and acquisitions which cause numerous infrastructure integration and interoperability issues. Needless to say, the ability to collect and present infrastructure data is challenging at best! Whichever organizations you speak with today, their top concerns on this subject are the ability to:

- quickly assess the impact of a service failure;
- identify the components that make up a service so that Service Level Objectives can be defined and availability measurements can be identified;
- assess the risk associated with the implementation of a change to the production environment;
- accurately cost and price their IT services;
- manage the financial lifecycle of their IT assets.

The IT process of Configuration Management and the related business process of Asset Management provide support for these business decisions, in concert with the other IT processes such as Change Management and Service Level Management.

How many times have you asked for a copy of the "configuration map" that is unique to your colleague's technical domain because you could not deliver on your infrastructure component operational responsibilities without knowledge of another area's infrastructure? The data contained in these "maps" is critical to the daily running of an IT infrastructure. Gone are the days when the resident shop "guru" is your single contact for information regarding your complete infrastructure architecture and all the associated components. The platforms that support service delivery are too complex and integrated for any one individual or group to identify. How many times have your Service Desk staff wished for a quick reference of all customers and their services affected by a failing IT infrastructure component? How often has your organization struggled with the identification of the technical components that enable the delivery of a service in order to set and measure meaningful service level objectives in support of Service Level Agreements? It is the underlying process of Configuration Management that positions an organization to identify, track, record and report on IT infrastructure data as it relates to the management, costing and delivery of IT services.

The Configuration Management process, as defined by the Information Technology Infrastructure Library (ITIL) best practice and detailed in the HP ITSM Reference Model, will be described in detail in this paper. Clarity will be provided around the confusion created by the terminology and practices of Asset Management, Inventory Management, Configuration Monitoring and Configuration Management. Integration between Configuration Management and its associated IT business processes will be also be discussed in detail. A variety of implementation success stories and strategies will be presented throughout this paper, covering the people, process and technology areas of focus that are essential to the successful activation of a Configuration Management process in an IT organization.

The intended audience for this paper is IT Management, Executives and Project Managers who are considering activation of a formal Configuration Management process within their organization and for IT professionals with a basic understanding of Configuration Management principles. This paper will highlight the requirements of the supporting tool set for the Configuration Management process, however, it is not intended to be a detailed technical discussion around the tools required to support a Configuration Management process activation in an IT organization. For more detailed technical information, please refer to the support documentation for the specific tools mentioned in this paper.

1.2 WHAT IS CONFIGURATION MANAGEMENT?

Configuration Management is a supporting IT process described in the ITIL best practice. Configuration Management is the process which manages and controls the revision of all components of the IT infrastructure and their current status relative to the production environment.

1.3 WHAT IS ITIL?

ITIL (The Information Technology Infrastructure Library) has been dominant as the only public best practice for IT Service Management in the European market for many years. ITIL defines the operational, tactical and strategic supporting and delivery processes required in the running of an IT shop.

Organizations choose to use the ITIL best practice to assist them with achieving control of their IT infrastructure and services while still maintaining the balance with the flexibility required to meet time to market demands. ITIL itself is a theoretical best practice for IT process management, focused on the quality support and delivery of IT services. However, theory is never quite as elegant when translated into an actual activation project. In fact, activation can become quite challenging and at times, even messy. Unfortunately, the translation of this theory into practice often results in failed projects.

It is interesting to note that organizations will often introduce new service and system management technology with the goal of improving the business. However, many of these implementation projects fail, as they haven't fundamentally altered their underlying business processes to accommodate the new way of doing business that the technology promised to deliver. In any major technology implementation, the integration between process, people and technology needs to be addressed. ITIL offers best practices that can be used to address all components of this "activation trio".

People, Process and Technology "Activation Trio"

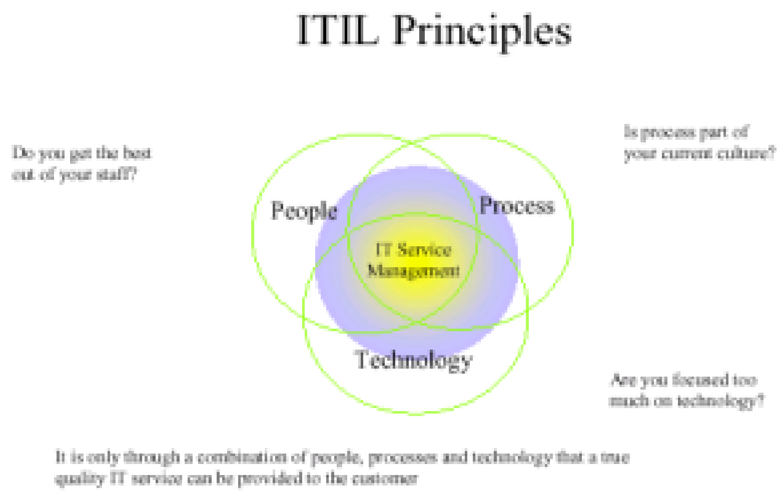

1.4 CONFIGURATION MANAGEMENT PROCESS IN DETAIL

Within the ITIL definition of the Configuration Management process, IT infrastructure components are referred to as Configuration Items (CIs). Configuration Management is the "documentation of the state of a CI" in the IT infrastructure. CIs managed by this process include hardware items, software components and object code, network items, documentation and any other element within the IT infrastructure that an organization needs to control. Configuration Management defines the relationship between CIs, such as parent-child, upstream-downstream and peer-to-peer.

Information about the CI is stored in an entity known as the Configuration Management Database (CMDB) that may consist of a single database or, in the case of larger or more complex IT organizations, the CMDB may be a single logical entity consisting of multiple distinct databases. Generally, the CMDB includes all documentation related to that CI. The Configuration Management process creates, maintains, tracks and reports on information that enhances the ability of other supporting processes to be effective, including the ITIL processes of Change Management, Incident Management, Problem Management, IT Service Continuity Management, Financial Management for IT Services and Service Level Management.

The process of Configuration Management addresses 5 key activities.

1. **Planning**: defining the scope, purpose, objectives, policies, roles and responsibilities and the process for Configuration Management.

2. **Identification:** specifying and identifying all components of an IT infrastructure and the design of the CMDB structure.

3. **Control:** the ability to agree and "freeze" CIs and then to make changes only with the agreement of appropriate named authorities (Change Management process) - controls the updating of CIs / CI status in the CMDB.

4. **Status Accounting:** the recording and reporting of all current and historical data concerned with each CI.

5. **Verification:** a recurrent series of reviews and audits to ensure that there is conformity between CIs and the authorized state of CIs as recorded in the Configuration Management Database (CMDB).

It is important for an organization to understand the full scope of activities they will need to address when considering the activation of a formal Configuration Management process. A closer look at each of the key process activities for Configuration Management reveals specific deliverables that will need to be addressed by an organization. Examples of specific deliverables are included in the following list:

- define levels of governance (delegation of authority);
- define configuration item naming and numbering conventions and label all CIs;
- define the relationships between CIs;
- determine the data migration process to the CMDB;
- define any automated discovery of CI information;
- define reconciliation policies and processes between configuration monitoring and configuration management;
- define the interfaces between multiple CMDBs, both for data entry and query (applicable if the CMDB is architected as multiple databases);
- determine level of CIs to be tracked at which level and the tracking policy;
- determine audit schedules / procedures;
- define CI registration process (includes identification of CI ownership – Service Delivery, Financial/Asset, and User associated to the CI);
- define the standard housekeeping procedures for the CMDB;
- maintain an accurate status of the CMDB - only authorized and registered components are active;
- interface with other support organizations to ensure the most effective use of the CMDB;
- specify Change Management report metrics in support of the Configuration Management Process;
- ensure other support services (e.g., Incident Management, Problem Management, Release to Production) are able to work effectively and efficiently with the CMDB;
- ensure all components relative to IT Service Levels are registered;
- produce and distribute standard and customized reports;
- ensure authorized procedures and work practices are followed;
- evaluate the Configuration Management process for continuous improvement.

As you can well appreciate, activation of a Configuration Management process in an organization is not a "quick-fix" undertaking. Configuration Management activation initiatives require a full understanding of scope, resource and financial commitments within the organization. Though the short-term cost and effort may appear high, the long-term benefits far outweigh the initial outlay. Let's take a closer look at the benefits.

1.5 WHAT ARE THE BENEFITS OF ACTIVATING CONFIGURATION MANAGEMENT?

Infrastructure data is essential to enable specific business processes in the organization. Your Service Desk agents, your technology service providers, your financial analysts, your account management teams and your business executives all require access to the data describing your IT infrastructure. Infrastructure data is also essential to enable the existing business processes in an organization, such as Asset Management, Provisioning and Procurement.

The key values to IT with the activation of a formal Configuration Management process are:

- **To Add Value to the Customer's Business While Reducing Costs:**
  "If you've got best practices in place for a well-managed environment, you can reduce your total cost of ownership by as much as 25% to 30%."

  Michael Gartenberg - Director of Research for Desktop Operating Systems at Gartner Group, Inc.

  This quote applies to the activation of IT Service Management processes as a whole in an organization, of which Configuration Management is one part. Initial process activation costs (project, personnel and technology) must be factored in to the overall long-term cost benefit. The costs associated with not implementing a formal Configuration Management process are much greater. If an organization does not have access to the infrastructure data required to efficiently manage incidents, they may not be able to avoid downtime penalties or production run losses and the organization will not have access to accurate, up-to-date infrastructure information on which to effectively base competitive business decisions.

- **Increased Stability of IT Environment:**
  Decisions will be based on timely, accurate infrastructure information. This assists the organization with avoiding mistakes that might lead to unanticipated downtime. With the requirement to identify and manage appropriate authorizations to make changes to the CMDB, security is enhanced and the risk of uncontrolled environment changes is reduced. Adherence to corporate infrastructure standards can be managed effectively.

- **Increased Service Quality**
  The provision of a facility for storing and retrieving documented customer expectations in the CMDB makes it easy to compare and improve actual service delivery.

- **Decreased Operational Cost**
  The organization will gain a faster, simpler, more thorough identification of infrastructure CIs, their attributes and their relationships. Process roles and responsibilities are clearly defined and understood, leading to a decrease in duplicated effort of resources.

- **Increased Flexibility**

Retrieved infrastructure data can be sorted, summarized, and reported on in whatever manner makes sense by whoever in the organization requires this information.

Other benefits, derived from the process objectives, include positioning the organization to:

- ensure all CIs are assigned to and are owned by the Configuration Management process;
- ensure that a complete record of all CIs and their configurations are captured;
- ensure that key relationships between standards and production systems are documented e.g. - PC Image V1.1 is on PCs # 1, 2,4 &10;
- ensure that CI records are under revision control and that a process is in place to update their current status;
- ensure effective communication within the organization on CI status;
- provide an adequate Configuration Management Database;
- create a baseline from which the complete IT infrastructure can be recreated after a disaster;
- report on metrics of the Configuration Management Process (i.e. process compliance and accuracy);
- ensure that all CIs can be tied to specific services and owners (Service Delivery, Financial / Asset and associated Users);
- provide the infrastructure data required to create business cases for justifying the development of new services.

1.6 ALL THIS SOUNDS GOOD TOO BE TO BE TRUE!

As with all positive ventures, there is always some risk associated with moving forward in a new direction. Listed below are the common barriers associated with the activation of a formal Configuration Management process:

- the discussion around the definition of CMDB scope, breadth, depth and overall structure could prevent the organization from moving forward as this tends to become overwhelming and organizations attempt to create the “perfect solution”;
- within smaller IT organizations, one supporting tool may be able to meet the requirements of the Configuration Management process, however, larger IT organizations may be sidetracked looking for or designing the “one tool” solution which may not be practical;
- lack of management commitment to sustain the initial development and activation project and to enforce Configuration Management policies and procedures on an ongoing basis;
- lack of clear responsibilities and adequate authority;
- confusion around which process or procedure to activate - asset management vs. configuration monitoring vs. configuration management vs. inventory management;
- confusion around the basic terminology of Configuration Management;
- inability or unwillingness to manage “willing compliance” of the staff with respect to the Configuration Management process procedures and policies;
- challenges around the definition and assignment of key process roles;
- the requirement to activate a Configuration Management process with a formal Change Management process – some organizations may not have an existing Change Management process in place;

- the size and scope of the definition of standards and naming conventions and the subsequent component labeling activity;
- the inherent complexity of the design of process governance levels and application within the organization;
- the complexity involved with definition of an audit and verification process;
- the lack of complete data collection and filtering tools to populate and maintain the CMDB;
- the inability to achieve the right balance between manual population, updating and housekeeping of CIs within the CMDB;
- the inability to develop clear and achievable Configuration Management metrics to manage and improve the efficiency and effectiveness of the process and the organization.

Though this list may appear lengthy, the fact that these barriers can be identified from the start positions the organization to develop effective strategies to eliminate these barriers throughout the activation project. Forewarned is forearmed!

1.7 BUT WE KEEP AN INVENTORY!

The dictionary describes an inventory as “a list of goods on hand, an itemized list of current assets”. Inventory Management is the process that produces this output. Most organizations schedule an inventory on an as-needed basis, either because of an audit or legal requirement or because of a specific IT activity that requires this information such as a product upgrade. The output of an inventory exercise is usually a flat file of data. There are no relationship descriptions as there are in a CMDB, and there are no procedures to keep this inventory up-to-date and accurate as there are in a formal Configuration Management process that is closely linked to a Change Management process. The Configuration Management process is mainly a continuous process while Inventory Management is rather a one-time activity run on an ad-hoc basis. The Inventory Management process is often seen as a sub-process of the organizations’ Asset Management process and is very often required by law.

Configuration Management is a process whose output provides the organization with an up-to-the moment status of each item in the infrastructure and the relationships between the various infrastructure components. This is achieved through its tight integration with a formal Change Management process within the organization. Each update to a CI record in the CMDB must be associated with a Change Management request (RFC – Request For Change). This link with Change Management is a key factor in the verification and audit of the CMDB data. When activated together with Change Management, Configuration Management is a building block for proactive Change Management. Without this link to Change Management, the configuration data is only as good to the organization as the last time it was updated.

Let’s take a look at our activities over the Y2K preparation timeframe. Most organizations paid a large sum of money for physical and logical inventories to be performed so they could accurately assess the impact of Y2K related bugs to the delivery of their services. And what was the value of this inventory output after completion of the data collection? Were processes in place to keep this data up-to-date whenever changes were made to the IT infrastructure? Were positive ROI’s achieved after the outlay of cash to complete this data collection exercise? Could the organization be assured of the ongoing accuracy of this data to support their key business processes and management decisions? Most often, the answer is no.

This is due to the lack of a process to keep this inventory data up-to-date after the initial collection. Even with an inventory file / database update procedure, the organization may not have effectively measured and reported on the process outputs or managed the enforcement of compliance to these metrics to ensure that high quality, accurate infrastructure data is maintained.

1.8 HOW IS CONFIGURATION MANAGEMENT DIFFERENT FROM ASSET MANAGEMENT?

The dictionary describes assets as "items on a balance sheet showing the book value of property owned". The Configuration Management process is often confused with the separate financial process of Asset Management. Asset Management is the business process in the organization that manages an asset throughout its financial lifecycle. The Asset Management process includes activities such as contract management, cost depreciation, total cost of ownership, total value of ownership and obsolescence for each asset. Although there are key links between these two processes, they are in fact distinct from each other.

The ITIL best practice defines asset management as "a recognized accountancy process that includes depreciation accounting"[1]. Asset Management systems maintain details on assets above a certain value, their business unit and their location. The relationships to other assets are not usually recorded. In some cases, organizations start with the activation of Asset Management and then move on to the activation of a formal Configuration Management process. The key point is to understand they are two distinct processes with unique activities. Typically, an inventory is included within Configuration Management. Once you have activated the Configuration Management process, you also have created an inventory. But the reverse is not true with Asset Management and Configuration Management. An organization may still require an Asset Management database separate from the CMDB, (depending on the CMDB scope set in the Configuration Management process) and will definitely require an Asset Management system, which delivers specific application functionality such as cost depreciation or total cost of ownership calculations and reports.

An asset may or may not be identified as a CI, based on the scope determination exercise in the planning and identification activities of the Configuration Management process. An asset can be purchased, received and placed in stock in the organization without being moved into the production IT environment. When the asset is defined as a CI and stored in a CMDB, the Configuration Management process activity of Status Accounting will need to define the relevant statuses associated with a CI throughout its lifecycle. When the asset is ready to be moved into the production environment, a Request For Change is prepared and entered into the Change Management process. Following the Build and Test and Release to Production process activities, the asset is released into the production environment and the status of the CI record is changed to the organization's defined production status.

The following diagram gives an example of the different statuses that could be associated with a Configuration Item (CI), when the asset is defined as a CI and stored in a CMDB:

[1] ***Best Practice for Service Support – ITIL The Key to Management IT Services*** – Published by the CCTA on behalf of the Controller of Her Majesty's Stationery Office - 2000

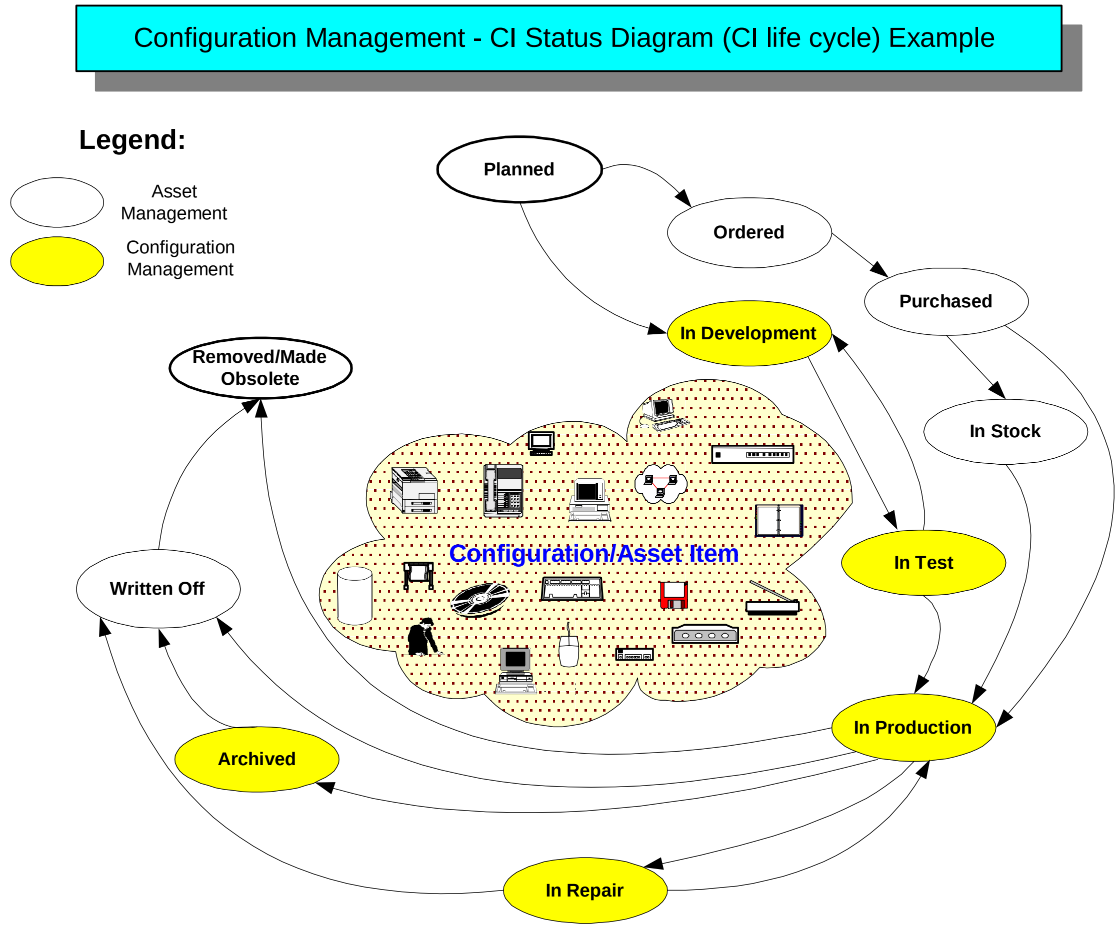

In the case where an Asset Management database exists separate from a CMDB, specific financial information about this asset would be stored in this database as part of a fully functional Asset Management system. The Asset Management system manages data input, field structures, standard financial and asset lifecycle (i.e. depreciation values) and specific asset management reporting.

If there is already an existing Asset Management database in the organization, (perhaps as part of a larger ERP implementation), the organization may choose to leave the asset database in place and connect to it from the Configuration Management process. But if there is not an existing asset database, do not create a new one, as the CMDB will contain enough information to act as asset database for the organization. However, the Asset Management system will reside separate from the CMDB. In this case, the asset from your CMDB can be exported as input to your Asset Management system.

1.9 I HAVE DISCOVERY TOOLS....WON'T THIS BE ENOUGH?

There is a difference between a Configuration Management process and the process activity of Configuration Monitoring. Configuration Monitoring is the activity of using discovery tools, such as

Microsoft Systems Management Server (SMS), HP OpenView Vantage Point Operations (VPO) or HP OpenView Network Node Manager (NNM) to automatically scan infrastructure components. Software agents supplied with the discovery tools would need to be loaded on each infrastructure component to be monitored. This scanned data can then be used to automatically populate a CMDB or to produce CI discrepancy reports, based on a match against the current data in the CMDB. An organization that has implemented a formal Configuration Management process would also use the auto-discovery function of Configuration Monitoring to complete the CMDB data verification activity as defined by ITIL. An Inventory Reconciliation sub-process then manages discrepancies between Configuration Monitoring and Configuration Management.

The organization may discover that discrepancies exist between the number of changes made to the IT environment and the corresponding number of change records recorded in the Change Management system. This would be an indication that the process performers are not applying the Control activity of the Configuration Management process.

Organizations that have not activated a Configuration Management process may choose to use the process of Configuration Monitoring to populate the CMDB with infrastructure changes that occur on a daily basis. They may even choose to not move beyond a Configuration Monitoring implementation.

> *One large industrial company understood that their culture and encouraged behaviors were not consistent with moving towards a formal Configuration Management process. They were not prepared to make the changes necessary to facilitate the move to a process which was fully integrated with Change Management, so they chose to implement an extremely sophisticated and expensive configuration monitoring environment. They determined the set of attributes required to manage the IT infrastructure, built filters to block the non-essential data elements and mapped these data elements to the CMDB fields. This strategy worked well for the organization by moving them to a position where they could reactively identify their infrastructure components. However, by stopping at Configuration Monitoring, this organization will always have infrastructure data information that is "one-day old", they will always be dependant on sophisticated monitoring tools and the ongoing expense of resourcing this skill set as opposed to encouraging new behaviors and accountabilities with their current staff. With Configuration Monitoring, there is also a greater need for skilled capacity analysis and management, given the additional load placed on the infrastructure due to the required software agents and subsequent polling intervals.*

1.10 SO WHAT COULD THIS ALL LOOK LIKE?

The following diagram illustrates a complete activation of a Configuration Management process within an organization. This diagram highlights the interfaces to IT Service Management support tools (change and problem management), links to the business processes of Asset Management, Facilities Management and Order Processing and the associated configuration monitoring and event handling tools:

Simplified model of a Configuration Management Process architecture as implemented at a large Bavarian financial service provider

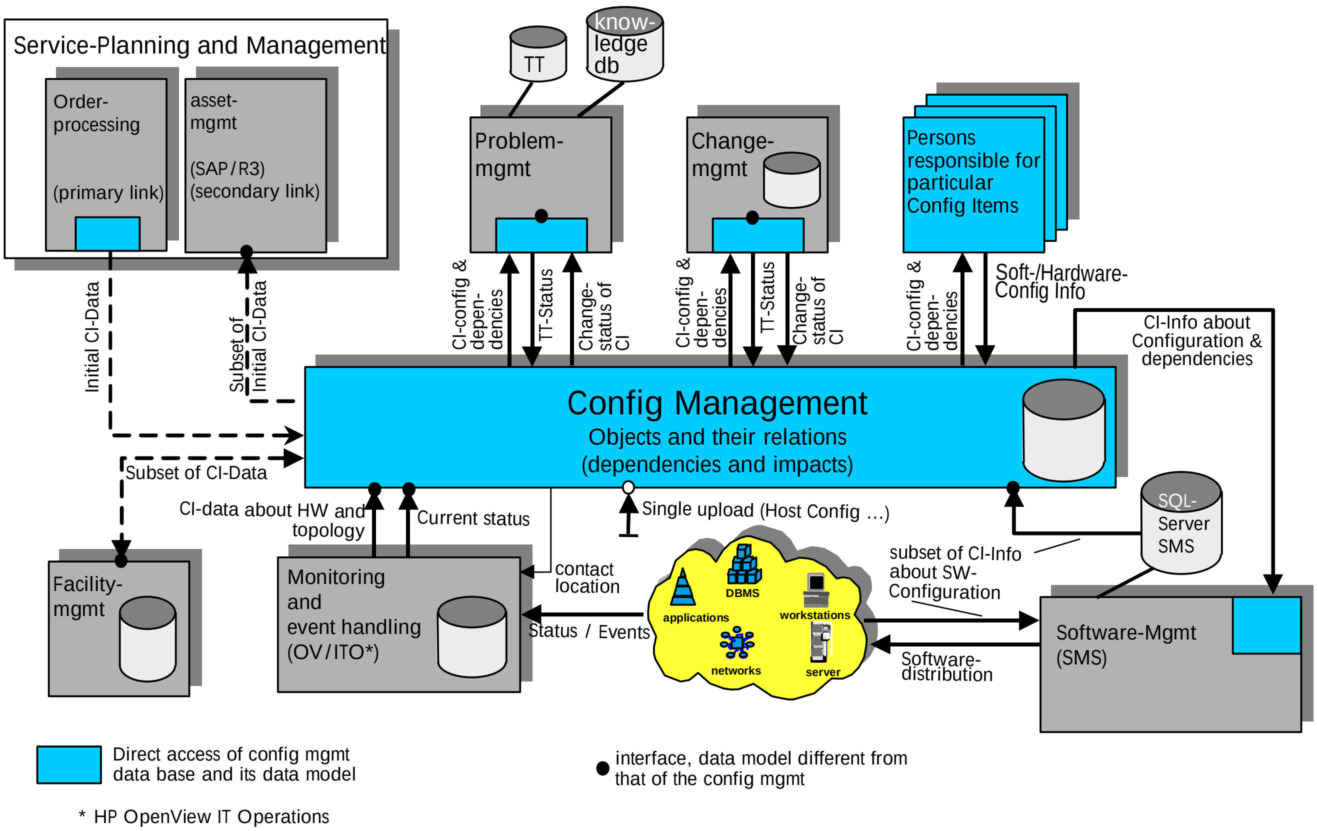

2.0 ACTIVATION STRATEGIES

2.1 HP IT SERVICE MANAGEMENT REFERENCE MODEL

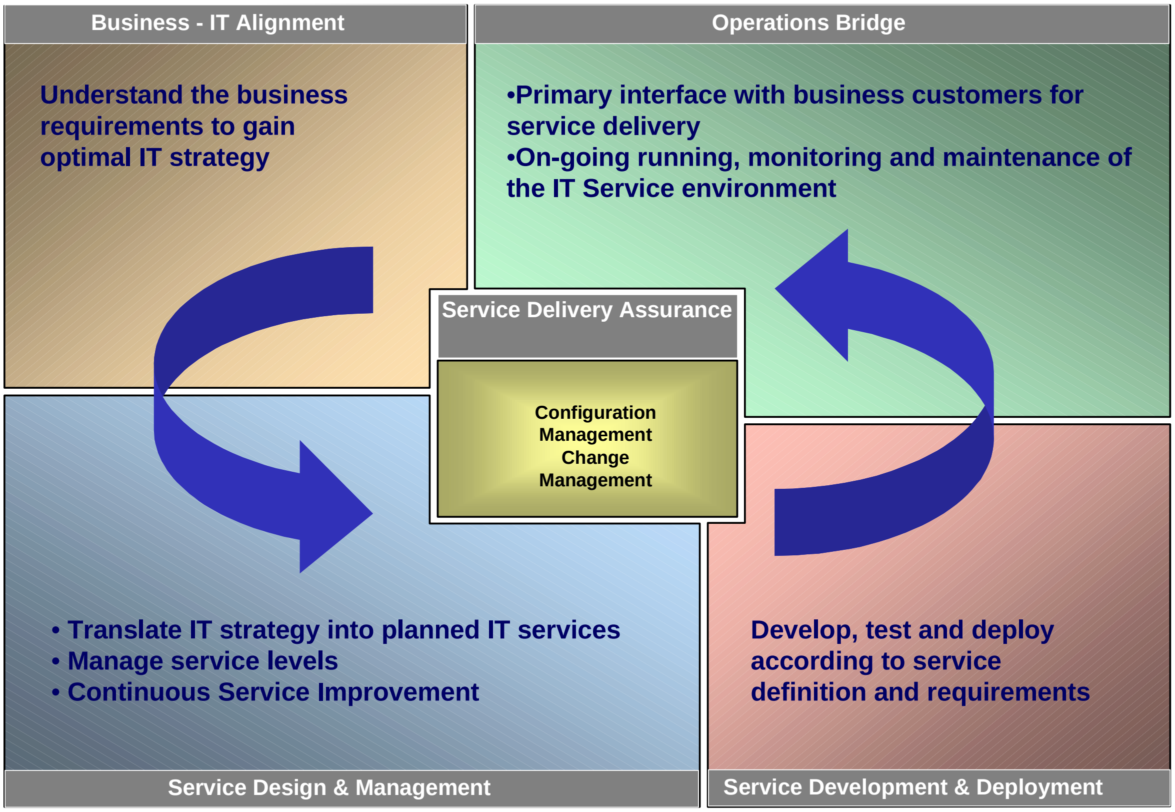

HP has developed an ITSM Reference Model based on ITIL as the underlying best practice. This model is divided into 4 quadrants: Business to IT Alignment, Service Design and Management, Service Development and Deployment and Operations Bridge all supported by the Service Delivery Assurance processes of Change and Configuration Management.

The benefits[2] of the HP ITSM Reference Model include the ability to:

- define current IT environment (as-is);
- identify gaps and desired state (to-be);
- use standard terminology;
- place ITIL into a framework to complete the picture of how IT as a business should be run;
- prioritize planned IT work efforts;
- identify critical process linkages;
- link problems to processes;
- link organization to services;
- target areas for potential process-enabling technologies;
- identify insourcing and outsourcing opportunities.

[2] HP ITSM Reference Model White Paper Version 2.0, Jeff Drake, 2000

The processes included in each quadrant are listed, along with their various activities, in the following diagram:

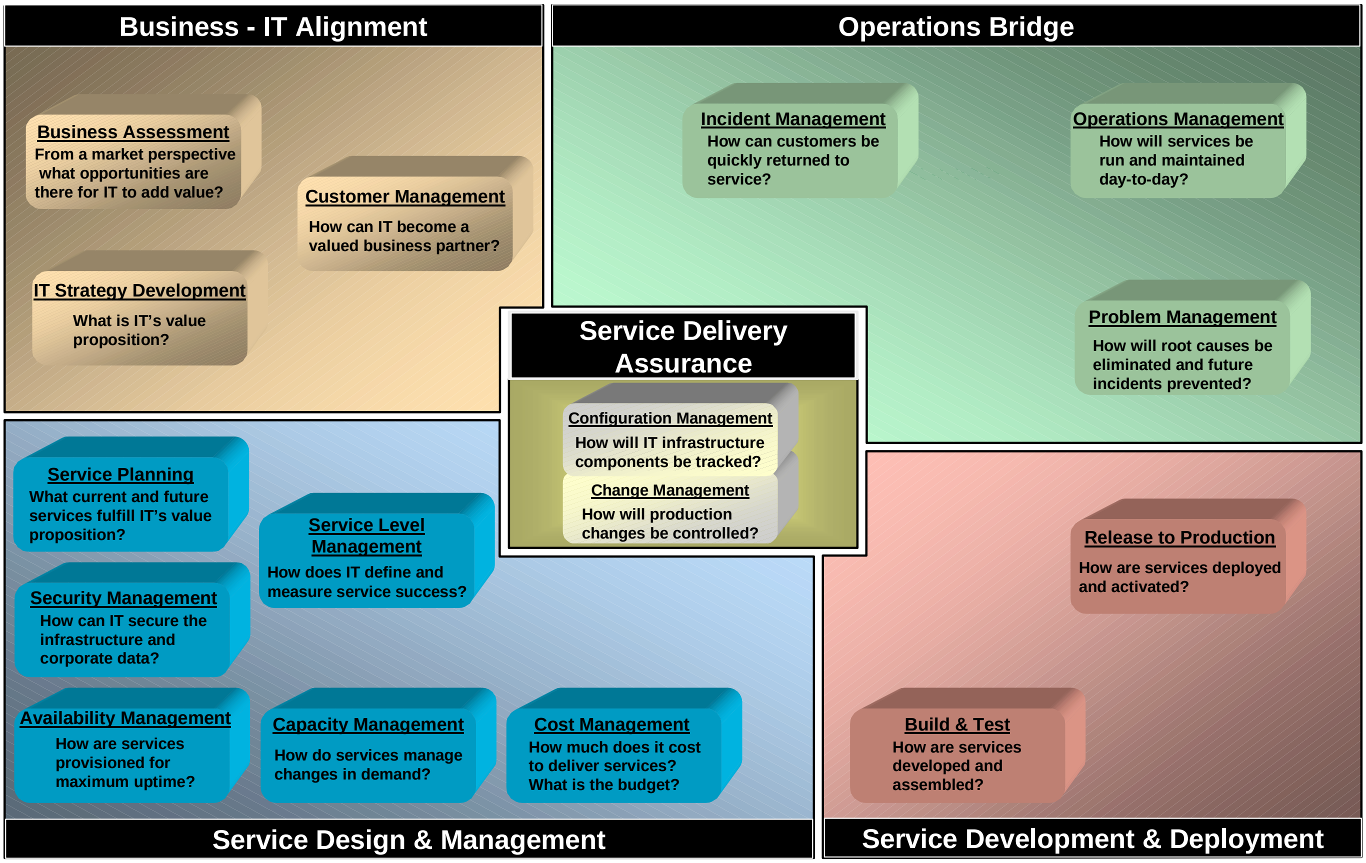

The HP IT Management Consulting Practice uses this reference model as a framework when activating the IT Service Management processes within an organization. The core processes, which anchor the ITSM processes, are the Change and Configuration Management processes. For more detail on the HP ITSM Reference Model, please refer to the January 2000 HP ITSM Reference Model White Paper Version 2.0, authored by Jeff Drake of Hewlett-Packard Consulting.

2.2 PROCESS INTERDEPENDENCIES

The following simplified diagram depicts the HP ITSM Reference Model process interdependencies with the Configuration Management process:

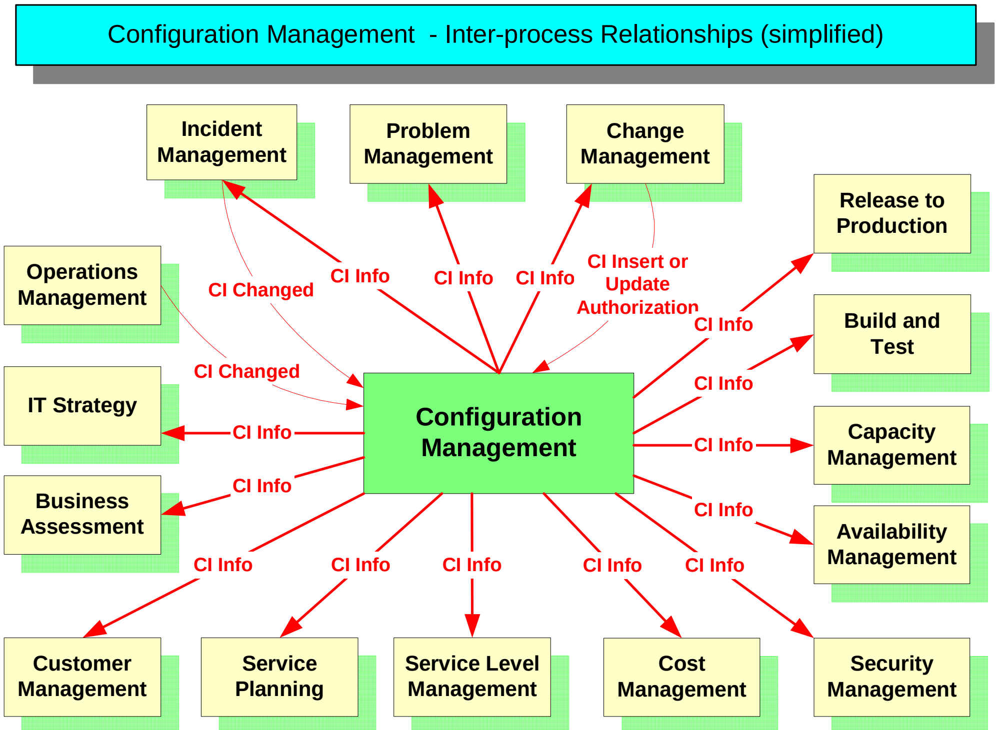

The interdependencies between the ITSM Reference Model processes and the Configuration Management process are described as follows:

Incident Management:

Configuration Management

- provides the infrastructure data required to assess customer impact of an IT infrastructure component failure
- identifies the CI-Owners for Service Delivery support, Financial / Asset ownership and associated User(s)
- data is used to correlate the CI with the appropriate SLA to determine the priority of actions and escalations

Incident Management

- provides a feedback to the Configuration Management process re CI accuracy

Problem Management:

Configuration Management

- provides the infrastructure data required to identify the root cause of incidents and to determine the course of action for problems

Operations Management:

Configuration Management

- provides updated information about infrastructure components in the production environment (e.g., production jobs, server configurations, IP address assignments, etc.)

- provides the infrastructure information required to perform day-to-day processing, to obtain status information and CI relationship information

Release to Production:

Configuration Management

- provides status changes for components, service functions, or end-to-end services; information about a production release
- provides the ability to understand and assess the impact of production testing on other services through the definition of CI relationships and dependencies

Build and Test:

Configuration Management

- CMDB is used to store document information about operating procedures, support procedures, recovery procedures, information about resources modified in the test environment or limited production environment, status changes for components, service functions, or end-to-end services and the registration of a master blueprint
- provides an understanding of the impact of component, service function and end-to-end service testing on other services and CI relationships, information required to prepare training materials and for documenting versions

Service Planning:

Configuration Management

- CMDB is used to store document information about external and internal design specifications for a service, service portfolio, service quality plan, service catalog
- data is used to determine the fit of existing components and service functions, with new / modified end-to-end services

Security Management:

Configuration Management

- CMDB is used to store information about security specific CIs, such as firewalls

Availability Management:

Configuration Management

- CMDB is used to store document information about availability designs, models, plans, subcontractor terms and relationships
- provides relationships between service components, service functions, and end-to-end services in order to design the availability reporting relationships
- provides a baseline snapshot of the IT infrastructure information which could then be used to assist with recreation of the IT environment in the event of a disaster or major interruption to service delivery

Cost Management:

Configuration Management

- identifies the IT infrastructure components that enable delivery of a service to the customer which allows for the accurate costing and pricing of IT services
- CMDB is used to store financial data documents such as service budgets and costing analysis reports

- infrastructure information can be used to evaluate service performance against a defined standard

Service Level Management:

Configuration Management

- identifies the IT infrastructure components that enable delivery of a service to the customer which allows for the accurate setting and measurement of Service Level Objectives in a Service Level Agreement (SLA)
- CMDB is used to store information about specific SLAs and registration of custom services
- data is used to prepare SLA exhibits and to prepare for service performance reviews

Capacity Management:

Configuration Management

- data is used to track CI performance characteristics which provides the current baseline relative to new demand and workload requirements
- stores document information about capacity design and plan, capacity analysis
- data is used to review relationships between service components, fitting end-to-end capacity and performance requirements specified in the design
- identifies the IT infrastructure components and their role within service delivery to measure and report that the particular component meets its capacity objective

IT Strategy Development:

Configuration Management

- data is used to identify funding limits for service investments
- stores document information for IT Strategy
- data is used to review summarized data about CIs to determine future directions and to determine skill categories (e.g., internet, NT, etc.)

Business Assessment:

Configuration Management

- data is used to compare existing services with market needs and to identify existing services

Customer Management:

Configuration Management

- is used to access current customer information and to update new customer information

Change Management:

Configuration Management

- data is used to assess the impact of proposed changes on IT services (including their supporting infrastructure)

Change Management

- provides the Control activity of the Configuration Management process (cannot complete an activation of a Configuration Management process in an organization without links to a formal Change Management process)

- must be activated first in an organization (before Configuration Management) in order to control and comply with the control activity
- provides authorization to change CI content / status

The Configuration Management process flows through the HP ITSM Reference Model as follows:

> *The IT Strategy process develops the overall architecture for the service. The Service Planning process develops the component (CI) requirements for the service after conferring with Availability, Capacity, and Cost Management processes to develop their respective component (CI) needs. The Service Planning process sends the service plan to the Change Management process as a Request For Change, at which time the new CIs are entered by the Configuration Management process into the CMDB and noted as status 'develop' upon authorization from the Change Management process. The Change Management process develops a project plan and issues a work-order to the Build & Test process for the new service. The Change Management process notifies the Incident, Problem, and Operations Management processes of the new service and its components being built. The Build & Test process notifies the Configuration Management process of the movement of CIs to a 'test' status. The Build & Test process develops any Operations Bridge process procedure documents, which can also be treated as CIs. The Release to Production process notifies the Configuration Management process of the change to an 'installed' status through the Change Management process. The Operations Management process accepts the new CIs and notifies the Configuration Management process of the change to a 'production' status through the Change Management process. The status update is reflected in the CMDB.*

2.3 PROCESS GOVERNANCE

In any government, policies are controlled at different levels. These levels reflect the domain of responsibility for implementing these policies. For example, a Federal level of government is not concerned with the details of how a Township level of government handles their garbage collection. The Federal level of government creates the policy, which states that garbage must be collected for the safety of the public and then delegates the authority and responsibility for implementation of this policy to the Township level of government. In many countries, we can identify at least two levels of governance; Federal and Township. Some countries may also be familiar with a third level of governance, often referred to as Provincial or State levels, which exists between the Federal and Township levels of government. For the sake of clarity in countries where these levels of government do not apply, we will use the terms Enterprise and Local to describe the Federal and Local levels of governance throughout this document.

For this discussion, the simplest way to view governance is "the delegation of authority and responsibility for actions". The Enterprise level policies take precedence over the Local levels. This allows for autonomy at the local level while still providing the required consistency across the whole. This discussion will focus on the most common implementation of governance, 2 levels - Enterprise and Local, as it applies to Configuration Management and CMDB structure.

Within an organization, the management and control of configuration items are typically controlled at different levels (i.e. the Enterprise level or the Local level). This enables the autonomy of the management and definition of CIs at the Local level, while still providing for consistency of CIs in the

Enterprise level CMDB. For instance, if too many CIs are controlled at the Enterprise level, the process will quickly become unmanageable and an unacceptable amount of bureaucracy will develop. Alternatively, if too much control is pushed to the Local level, this may result in a state of "anarchy", where the organization no longer has standard naming conventions or data storage methodologies, resulting in a loss of the big-picture view of the corporate IT infrastructure. Hence, there is a need to separate and define the scope of the CIs that must be managed at the Enterprise level and the CIs that can be managed at a Local level.

Using the category "Servers" as an example, a Enterprise level CI would be an Application Server running a mission-critical application and a Local level CI would be the File/Print Server within a technology domain in the organization.

An organization would be wise to apply this concept of governance to its CMDB design. In smaller organizations, or where there is a central help desk within the organization, it may make more sense to create one single level of CMDB. In the case of a central help desk, you could have a single, central CMDB and may use the distributed local IT-personnel for verification purposes. In larger organizations, the 2 levels of governance may be applied to the CMDB design. In this case, we recommend this governance strategy be used by organizations when beginning a Configuration Management process activation in a large and complex environment.

The Enterprise level CMDB would be maintained by an Enterprise Configuration Manager so that basic information regarding the Enterprise level CIs will be made available to the organization, such as an audit history of changes against CIs or basic CI relationships. Very limited information should be maintained at the Enterprise level CMDB. For example, a simple Enterprise level CMDB may contain:

- CI Name
- CI #
- CI Owner
- incident dispatch group contact information
- CI Version #
- limited software / hardware info

The Enterprise Configuration Management process manages and controls the revision to CIs that have been released to production at an Enterprise level. An easy way to determine the level of governance for a CI is to assess if a change to a CI at this level will have an impact one or more services delivered to the customer. If a change to the CI has no direct impact on other CIs required to enable the delivery of service to a customer, then the CI can be managed at the Local level. Otherwise, the CI must be defined and maintained at the Enterprise level. In the case of an Enterprise level CI, a corresponding Enterprise level Change Management process is often activated in an organization and followed for all changes to this CI, to ensure the impact of the change is properly assessed and managed.

In the case of two levels of governance, the Local level Configuration Management process manages and controls the revision to CIs that have been released to production at a Local level. The CIs managed at this level are locally-contained and a change to a CI at this level usually will only affect that particular "workgroup". If a change to a Local level CI affects the functioning of components required to enable the delivery of corporate services, then the CI needs to be redefined as a Enterprise level CI. Local level Configuration Management support tools (i.e. CMDB) are selected at this level to optimize the process of managing the configuration and monitoring of a CI. For example, the router support group may use a product such as HP OpenView Network Node Manager (NNM). Local level CMDBs are typically technology specific and are built to manage the configuration and monitoring of one specific platform, such as NT or UNIX.

It is rare to find an organization that requires 3 levels of governance. In the case where a third level of governance is required, this third governing level could exist along each of the Lines of Business and would hierarchically fit between the Enterprise and the Local levels of governance. Subsequently, a Local level Configuration Management process would then exist within each Line of Business for each technology unit or functional group. A CMDB unique to this level of governance would be maintained by a new Configuration Manager and this Configuration Manager would also ensure that their CMDB is linked to both the Enterprise level CMDB and the associated Local level CMDBs. This is a complex implementation, due to the degree of governance employed however, this model of governance is unavoidable in some organizations. Once again, organizations are better suited to apply levels of governance that will logically align with the way the company itself is structured.

An example of a 3 level governance model was implemented in a large telecommunications company. This company had outsourced all of its IT functionality to many different service providers; one provider for wide area network provisioning and support, one for application development and support and one provider for legacy system support. The Enterprise level CMDB included CIs that were integrated between the service providers to enable service delivery to the telecommunications company. The second level CMDBs were the databases maintained within the individual service provider organizations and the Local level CMDBs were defined as the individual technology domains within a single service provider company. Complex, yet effective for the telecommunications company to manage both configuration data and changes to the IT infrastructure across the multiple service providers.

Whatever levels of governance are applied within an organization, it is important to note that a CMDB or a Configuration Management process does not replace the expertise regarding IT infrastructure components and their inter-relationships that exist within an organization. The CMDB is meant to augment the collective knowledge and expertise of the organization by providing the supporting process and enabling technology to share the information across the company.

2.4 ENABLING TECHNOLOGY

The Configuration Management process relies heavily on supporting technology. Millions of dollars have been spent on tools to enable this process within organizations. The following chart describes the characteristics of the tools (single tool or multiple tools) required to support the Configuration Management process activities in an IT organization.

	Configuration Management Process - Supporting Tool Abilities
Security:	security controls and interfaces that control access to configuration information in the Configuration Management Database (CMDB) on a need-to-know basis
	Integration with a Change Management system such that updates to a CI record are not allowed without the corresponding approved Request for Change (RFC) record status
	Management of the environment from a Web browser
Usability:	Add new CIs, modify existing CIs, and remove obsolete CIs from the CMDB
	Specifically track CI version, status, owner, and relationships to other CIs
	Run standard and ad-hoc queries both offline and in real time
	Maintain historic / audit trail information for CIs
	Interface with configuration monitoring tools to initially populate the CMDB
	Accommodate users spread across the enterprise in multiple geographic locations
	Compose or decompose complex CIs
	System maintenance of referential integrity
	Management and administration of system from multiple geographic locations
	Robust solution (high performance, low maintenance) thereby requiring minimal support resources / effort
	Support CIs of varying complexity from entire systems to single hardware and software items
	Provide graphical representation of CMDB sections such as CI relationships or network diagrams
	Perform automatic validation of input data (for example, CI name is unique)
	Support the management and use of CI baselines, including support for reverting back to a trusted version
	Automatic generation of CI relationships when new CIs are added
Scalability:	Easy scalability to accommodate the steady growth of new CIs
	Manage performance impacts and capabilities as the # of CIs grow by orders of magnitude
Integration:	Integration with other IT Service Management support tools such as incident, problem, change, service level management products
	Integration with a systems management framework, and what frameworks are supported (i.e. Microsoft SMS, HP OpenView VPO and NNM)
	Exchange of data across related systems. (i.e. data from a network management or problem management system being fed into a Configuration Management system)
	Easy programmability of the solution's interface(s) to systems management or IT Service Management tools
Auditing:	Report generation on CI inventory to facilitate configuration audits
Reporting:	Easy generation of standard and ad-hoc reports - reports run efficiently in real time
	Interfaces to available report programming technologies

	Provide import/export capabilities
Miscellaneous:	Asset Management functionality - ability to manage the financial aspects of assets (e.g. depreciation, obsolescence, maintenance, chargebacks, budget forecasting, licensing)
	Open architecture supported vs. proprietary components in the proposed solution

Within a small IT organization, a single supporting tool can often fulfill the requirements of the Configuration Management and Asset Management processes and a single CMDB can be created to store all the infrastructure data. However, the “single-support tool” option may not be practical or feasible for larger IT organizations.

In the case of large IT organizations, the activities involved to perform market research, (in order to identify the “one tool solution”), or with an attempt to build a single, powerful CMDB that can manage complex and detailed relationships, including automated updates and linkages with Change Management tools, can become overwhelming. In fact, many organizations get no further than this stage in the Configuration Management process activation planning process. Because of this complexity, it is important to adopt the strategy to begin CMDB development as simply as possible. Using the concept of governance, an Enterprise level CMDB can be produced simply by using a spreadsheet application and a low maintenance database front-end tool. Minimal CI data is required to begin the implementation of an Enterprise level CMDB. Once process maturity is achieved at this level, more sophisticated integration with the various levels of CMDBs can be implemented.

2.5 DEFINITION OF CMDB SCOPE AND DEPTH

When planning to activate a Configuration Management process, the organization will need to determine the breadth, depth and scope of the CMDB and the associated costs linked to these decisions. As detailed in the following diagram, the CMDB breadth includes the definition of CI attributes, the CMDB depth or level includes the number and degree of CI relationships that are maintained and the CMDB scope involves identifying which CIs will be stored. The deeper and broader the scope and level of the CMDB design exponentially increases the cost of initial population and ongoing maintenance of the infrastructure data.

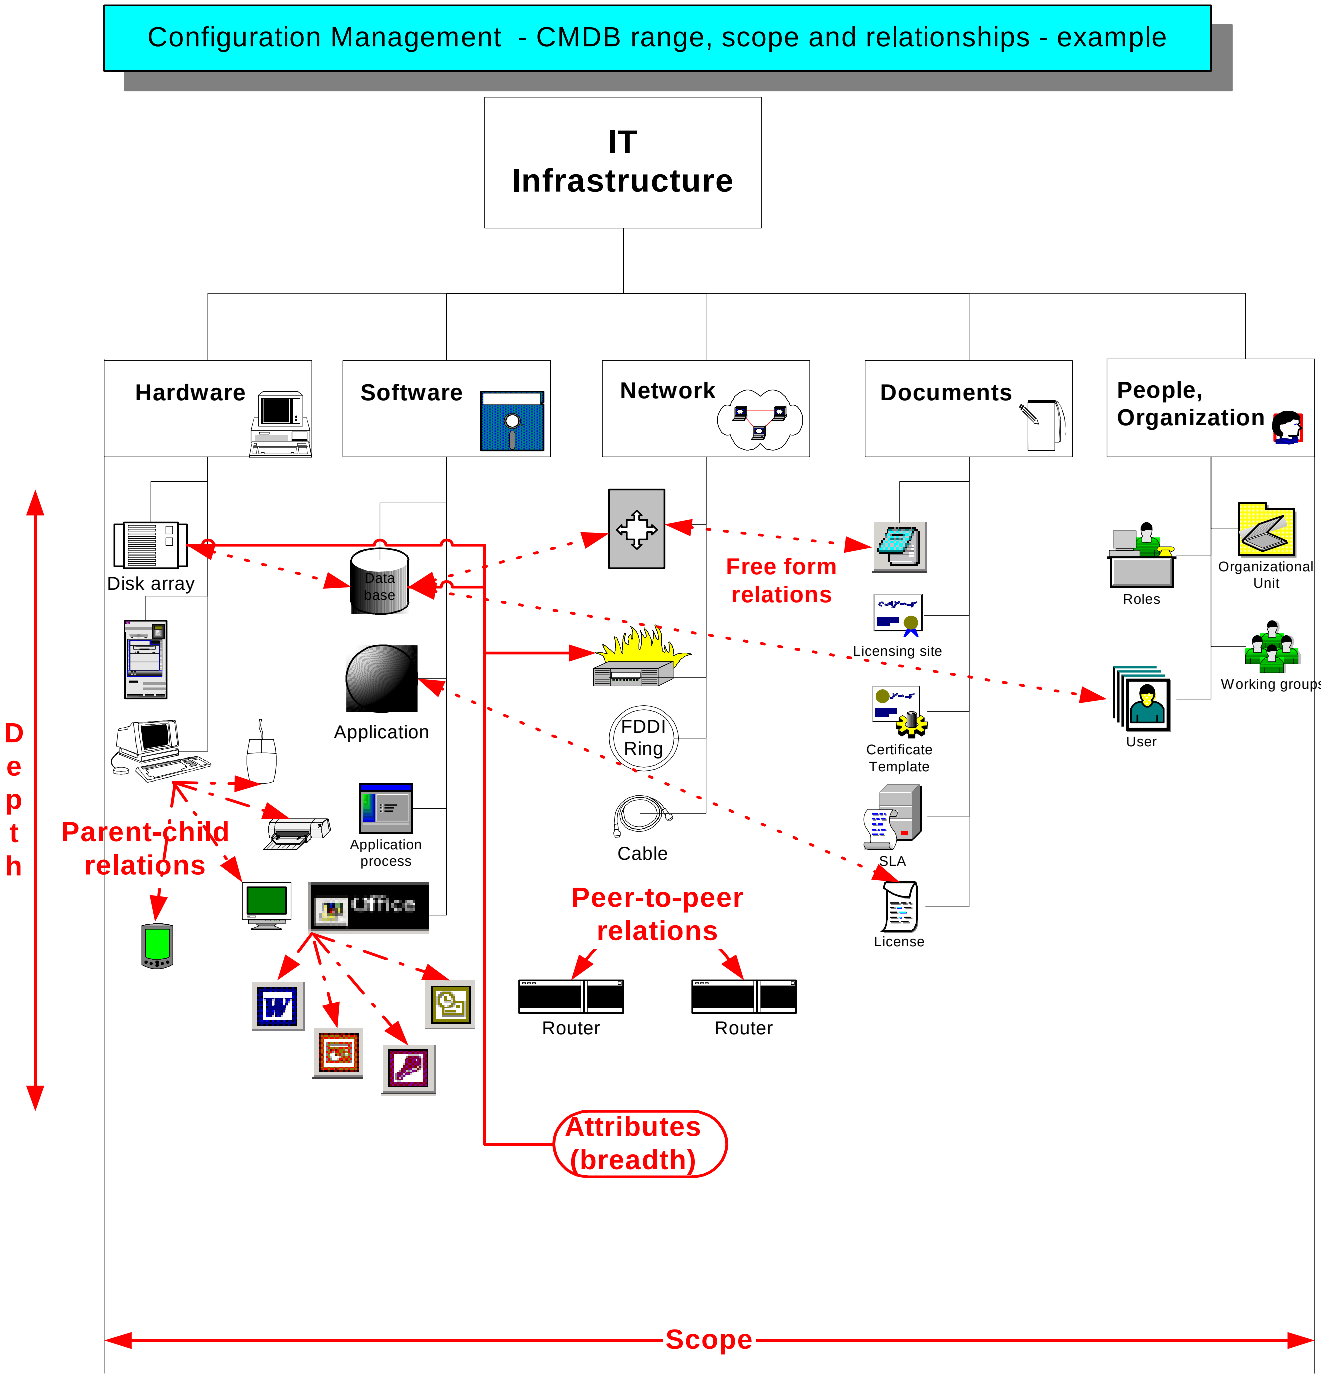

Whether you are a large or small IT organization, the rule to follow is "Keep it Simple"! Continually ask yourself: do you need to identify each component of the PC, such as the mouse or the power cord, or if you can track the CI at the PC level? Who is going to use this information and how are they going to use it? If it is not directly relevant to the ongoing delivery of service, or it can easily be queried on or stored elsewhere, then perhaps it does not need to be stored in the CMDB to this detail. As all changes to the CIs stored in the CMDB will need to managed through an administrative Change Management process, a good "rule of thumb" to follow is to only define CIs that need to be managed through a change control process.

> *A large European public sector organization spent a great deal of time debating the question of whether the wall outlets should be part of the CMDB or not. As a compromise, the structure was added, but until now (three years later) it is still not used. It was recommended to the customer that trying to keep track of the wall outlets by means of an administrative change process was not feasible and that loading wall outlets should wait until they had selected a tool specific for administering lines and wall outlets. When the specialized tool was in place, that tool could then be used as a feeder into the CMDB.*

Apply the concepts of governance to aid with initially managing the scope of your CMDB definition. In a large IT organization, begin with the definition and implementation of an Enterprise level CMDB. Immediate gains are experienced with provision of a single access to corporate infrastructure data. The challenges associated with CI-Ownership identification (Service Delivery ownership, Financial / Asset ownership and associated Users) is encountered on a smaller scope and the organization has time to develop effective strategies in support of further process penetration throughout the company.

> *A large utility company took this advice to heart. They collected all the "technology maps" that existed in each of the technology organization units. They identified each group that created these maps and asked if they would be willing to store these in an Enterprise level CMDB for all in the organization to access. The manager of each technology domain agreed not only to store these maps in one place, but to keep them up-to-date when changes were made in the environment. The Configuration Manager built a very simple spreadsheet application and linked the CI record to a visual display of the corresponding maps. In time, they built a front-end application entry point for presentation, data entry and reporting. The delivery organization was delighted to have access to this information in one application. Even this small step was a huge improvement over what was in place!*

If multiple, logical CMDBs are designed, single primary keys will need to be defined to link to the Local level CMDBs, such as CI # or CI (logical) name. Access to the Local level CMDBs must be provided to the people requiring more information than is contained in the Enterprise level CMDB in order to manage service delivery. Don't recreate what other technology companies have already spent time perfecting for their specific platforms. HP OpenView Network Node Manager was developed to support a complex network environment; Microsoft SMS is used to support an NT server environment. Don't get caught up in recreating this specialized functionality in your ITSM tool if they already exist in your company. Of course, if you are a start-up site or small IT shop, by all means use the full-functionality of the Configuration Management engines that exist in your IT Service Management support tool.

> *As an example of what can go wrong with CMDB design and development, an internal consultant in an IT organization once spent three full years trying to recreate the functionality of a network monitoring and management tool in a newly purchased ITSM tool. Three levels of relationships were designed and implemented to show:*
>
> - *the relationships between routers (peer-to-peer and parent child);*
> - *which port on the upstream router connected to which port on the downstream router;*
> - *identification of the telecommunications circuit connecting the ports and routers.*

This was a highly complex undertaking in order to recreate a function that the network management product was more than capable of performing. The reason for this was that the company was experiencing compliance issues to ensure the network management tool was updated when changes were made. It had nothing to do with the technology! In this case, the network management product should have been defined as a Local level CMDB and the appropriate process metrics and reporting put in place to manage compliance within the Local organizational units.

2.6 NOTES ON BUILDING A PHYSICAL CMDB, CI ATTRIBUTES AND DISCOVERY TOOLS

Many organizations try to import every data item that an auto-discovery / collection tool can provide into their CMDB. When physically mapping all of this information into a CMDB tool, the technicians will need to create many custom fields to accommodate the imported attributes. Just because a discovery tool can obtain this level of information, it may not necessarily add value to the organization's CMDB. For example, the packaging group uses Microsoft SMS to release software to the customers. If the packaging group requires the version of *DLLs* on a specific platform, Microsoft SMS has the ability to query for that information. The packaging group already has access to this data through the tool. What value would there be at the Help Desk in having the version of *DLLs* on a customer's PC for logging service calls and incidents?

This leads into the notion of combining your CMDB tool and your discovery tool(s) to create Enterprise level and virtual CMDBs. Keep the minimal information required in the Enterprise CMDB and if more detail on a specific CI is required, then access the "virtual" CMDB to get this information.

Key Points on Physical CMDB Design:

1. Keep the number of attributes to a minimum. Have the Incident Management group provide you with the minimum attributes they would require to perform the Incident Management process (and likewise for each ITSM process). It is easy to add more attributes at a later date. Use the analogy "If you were on a desert island and could only have 4-5 attributes, what would they be….?"
2. Ensure that your discovery tool and CMDB structure match. If your collection tools support three-tier data gathering, your CMDB tool should also reflect the same structure. You do not want to spend hours of data mapping time when you update or refresh your CMDB.
3. Ensure that your discovery tool and CMDB classifications match. If Microsoft SMS uses a specific classification for a shrink-wrapped application, make sure it is the same in your CMDB. This is a real time saver when mapping data fields.
4. Template your CIs. You can make the assumption that all PCs have a mouse and a keyboard, so create a PC template and incorporate these components into it. For "like" CIs, you can also template hard disk size, installed memory, NIC and video cards.
5. Build meaningful relationships between CIs. CMDBs with little or no relationships are cumbersome to search and perform incident analysis and problem management. Ensure that you can make the relationship between a person and their PC, and the

same for servers and their installed applications. It is difficult to break the relationships between CIs if you decide to reduce the depth of your CMDB at a later date.

6. Ensure that most of the data in the CMDB can be automatically updated from your discovery tools. Audit your moves, adds and changes (Change Management process) to ensure that the CMDB is being updated before the closure of the service request.
7. Design your CMDB from a "Service" view. Ensure that only business critical CIs are entered into the CMDB and that all the components of a "Service" view are accounted for (i.e. networks, firewalls, servers, and applications).
8. Use "Type" CIs whenever possible. Normally, for each occurrence of a CI in the infrastructure, a CI record in the CMDB is created. For an application such as an office automation software package, you need only record that CI once, but ensure that you have multiple instances of it. Unique CIs are necessary when you want to track specific attributes such as serial number and asset tag information.

3.0 PEOPLE CONSIDERATIONS – MANAGEMENT OF CHANGE CHALLENGES

So far, this paper has discussed the process and technology aspects of activating a Configuration Management process in an IT organization. The final point for discussion in the process "activation trio" is the "people" considerations. This concept is one of the simplest to understand yet is the most complex to manage and apply throughout the process activation project. This section addresses some concepts and approaches to success process activation.

3.1 KEY PROCESS ROLES

The table below lists the key process roles in a Configuration Management Process:

Process Role	Accountabilities
Configuration Management Process Owner	Owns the Corporate Configuration Management Process (at all levels).
Configuration Manager	Manages and co-ordinates all activities necessary to identify, control, track and audit all CIs – could be Enterprise, Domain or Local Level.
Configuration Coordinator	Records and maintains all CIs in the CMDB. Maintains the configuration CI dependency matrix. Maintains the CMDB. Maintains the Configuration Management System for data entry and reporting.
Service Delivery CI-Owner	Accountable for the performance, reliability and availability of the CIs they own. Ensures accuracy of CMDB data for owned configuration items. Ensures modifications to owned CI data are made as part of the Change Management / Release to Production processes.
Asset Management CI-Owner	Owns the financial accountability for the CI (asset).
User CI-Owner	Is associated to the CI through relationship definitions (i.e. this User is associated with this PC).

One key process role is that of Service Delivery CI-Owner. These are the people who are accountable for the performance, reliability and availability of the CIs they own. They are measured on meeting the Service Level targets within their technology delivery platforms. They are the key stakeholders in ensuring that configuration management information is available and accurate. They are the people who create the "configuration maps" for their systems that are essential to other organizational domains in order to deliver service across multiple infrastructure components. It is essential that this information is available in order to build, design and run the client's services.

The mapping of the Service Delivery CI-Owner process role usually matches perfectly to an organizational unit that is responsible for administering changes to the IT infrastructure. Typically, the Service Delivery CI-Owner is the manager of the technical domain within the organization. The Financial / Asset CI-Owner is identified for use in the Asset Management process for cost / budget purposes. The User CI-Owner is used identified in the CMDB for use primarily by the Incident Management process to link a CI to a person.

Activation of a Configuration Management process does not necessarily consist of building a single new organizational "group" that is responsible for the data entry updates to the various configuration management databases for the organization. In large IT shops, this would be impossible, as a single group would not have the required technical knowledge and experience to correctly portray the infrastructure topologies. Instead, each Service Delivery CI-Owner is made accountable for the accuracy and currency of their CIs in each level of CMDB across the organization. These updates are controlled and authorized through the activities within the Change Management process. It is the responsibility of the Change Management process to make sure that infrastructure alterations are planned and communicated to the Configuration Management process. However, the Service Delivery CI-Owner is the individual or group of individuals who have the authority to update the CI data in the CMDB as part of the Control activity of the Configuration Management process.

The organization will need to map these new process roles to the existing organization. Most processes involve many "organizational silos". The following diagram is a sample flow of a process across the organizational "silos":

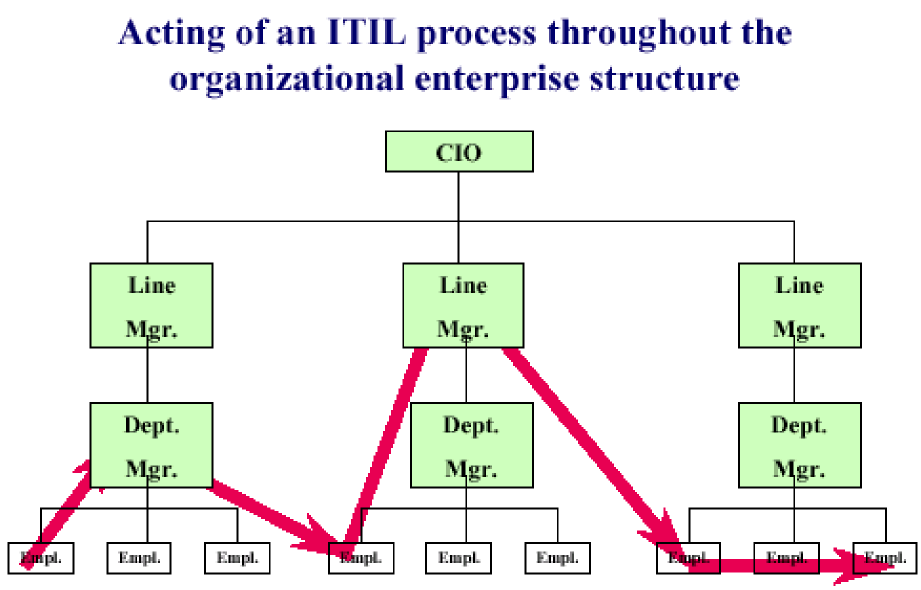

In this matrix-management environment, the process activation project team must clarify and communicate the difference between an organizational function and a process role. A function in an organization describes the job with a job title, the job activities and the direct reporting structure. For example, your job title may be LAN Administrator reporting to the LAN Manager. A process role describes the activities you perform in the specific ITSM process. You may perform the role of Service Provider in the Incident Management process, Change Initiator in the Change Management process and Service Delivery CI-Owner in the Configuration Management process. Each of these process activities then roll-up to become the activities listed in your job description.

It is important to have representation from your Human Resources (HR) department on the Configuration Management process activation team. Although it is often less of a challenge with the activation of a Configuration Management process, (as opposed to the activation of other processes such as Change Management or Incident Management), HR issues may arise around assignment of process roles. Gaps may be identified, where the process activities are not performed at all, or overlaps may be noted, where many groups deliver the same activities, resulting in redundancy and high cost of effort to the organization. Some IT shops may also have union representation. In this case, it is imperative to include the appropriate union members in the start up of any process activation project in order to ensure a successful project outcome. Planned and executed communication strategies are an organization's key tools to manage the human responses to change experienced with process activation projects.

A key benefit of the ITIL IT Service Management processes is that they describe and formalize what is required to run an IT organization. They are based on common sense and most of the activities identified in these processes are already performed within the current IT organization. This is not a re-invention of the wheel. The benefit of adopting these processes include:

- the identification of redundant activities resulting in the information required by management to achieve a reduction in duplicated effort;
- the ability to apply and communicate, in a common language across the organization, the process roles and responsibilities.

3.2 MANAGING "WILLING COMPLIANCE"

The biggest challenge to the successful activation of a Configuration Management process is the management of "willing compliance". The Service Delivery CI-Owners must willingly update the CIs they own in the CMDB (this includes each level of CMDB for many-layered governance implementations). There are various methods for reporting on and encouraging the key behaviors required to keep an organization's CMDB up-to-date when changes occur in the IT environment.

One way to understand the "right" way to do things is to study failed attempts. An example of mismanaged willing compliance was observed in a large IT organization. The organization had made a significant investment in a network monitoring and management platform. The ability to auto-discover and build sophisticated network service maps was inherent with this tool, however, certain organization specific information needed to be manually updated whenever there was a change to the environment. The organization had defined an update process step as a result of any changes, and metrics were produced and distributed to report on exceptions in the process. Even with all these steps, the organization was not able to influence "willing compliance" to keep this information up to date.

Dave Packard, the co-founder of Hewlett-Packard, once stated, "show me how a person is measured and I'll tell you how they behave". Had this organization identified CMDB update compliance as a business fundamental, with consequences for non-compliance linked to pay-for-performance or promotion incentives, the focus could be shifted towards meeting the metric targets.

Instead of managing the change in behavior through performance incentive initiatives, the organization decided that better technology was the answer. They purchased a high-end IT

Service Management Tool with Configuration Management / Asset Management modules. They automated the link between configuration updates and their change management process. However, the required level and detail of automation was so sophisticated, they had troubles training and retaining the staff to support the application once the contractors had left. The people using the tool had no understanding of how to back out the automated updates that were made in the CMDB when a change was backed out because the technology was too complex, resulting in data that was no longer accurate. This company is now adopting the HP ITSM Reference Model, based on ITIL best-practices, to address how business needs to be conducted within an IT shop, using their existing tools.

4.0 CLOSING

In closing, organizations that operate using a formal Configuration Management process will be successful in the support of the essential business processes involved in running an IT shop. Accountability of the individuals to ensure information on the IT components is up-to-date and accurate will be clearly defined, communicated and measured. Supporting tools will be in place to facilitate CI identification and CMDB infrastructure data verification for accuracy. Other IT support and delivery process, such as Incident Management, Change Management and Service Level Management, will have the access to the information required to make both the common and the complex business decisions essential to running an IT organization. Business units will have access to the information required to support the enabling of their applications through IT, such as a provisioning process in a telecom company or a procurement process for ordering new IT equipment.

This paper has identified various strategies for the successful activation of a Configuration Management process in an IT organization by paying close attention to the people, process and technology "activation trio" set of activities. Remember, keep it simple, start slowly and begin to immediately realize the benefits of successfully activating a Configuration Management process in your organization.

Reference Material:

HP ITSM Reference Model White Paper – Version 2.0, Hewlett-Packard Company, Jeff Drake, 2000

Configuration Management Process Guide – Version 2.0, Hewlett-Packard Consulting (Canada), 2000

Best Practice for Service Support – ITIL The Key to Management IT Services – Published by the CCTA on behalf of the Controller of Her Majesty's Stationery Office - 2000

