

From a technology-oriented to a service-oriented approach to IT management

A. J. Keel
M. A. Orr
R. R. Hernandez
E. A. Patrocinio
J. Bouchard

In this paper we focus on information technology service management (ITSM), a new direction for managing information technology (IT) organizations that requires a change from a technology-oriented to a service-oriented approach to IT management. We first examine the main challenges to implementing ITSM, which span four broad areas: process, people, technology, and data. Then we consider two strategies for implementing ITSM: top down and bottom up. Finally, we illustrate how a customer can implement ITSM by describing a simple scenario in which a business implements ITSM by using IBM Service Management solutions and tools and the top-down approach.

INTRODUCTION

Information technology service management (ITSM) is a discipline for managing organizations providing information technology (IT) services from a customer's perspective. The customer perspective implies a shift from a technology-oriented to a service-oriented approach to IT management. This shift changes the way that IT and business work together to define and deliver IT services.¹

The growing dependency of companies on IT, whose infrastructure has grown increasingly complex, coupled with the demands to reduce costs and comply with new regulations, has driven companies to search for ways to increase the efficiency of their IT operations. ITSM is the new direction that the business world is adopting.

Many business are also adopting the set of best practices known as the Information Technology

Infrastructure Library** (ITIL**), which provides a process-based framework for implementing ITSM.² However, even though many IT organizations are familiar with, or even well-educated in ITSM and ITIL precepts, many challenges remain.

The paper is organized as follows. First, we describe the challenges that IT providers face when implementing ITSM. Then, we discuss two approaches to ITSM adoption: top down and bottom up. Then we show a simple example which illustrates how a customer can implement ITSM using the IBM Service Management solution set and the top-down approach.

©Copyright 2007 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the Journal reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to republish any other portion of the paper must be obtained from the Editor. 0018-8670/07/\$5.00 © 2007 IBM

CHALLENGES TO ADOPTING ITSM

Traditional IT departments have grown into resource- or function-specific organizations. Historically, this resource-oriented organizational structure is a result of new support teams being created as new types of IT resources are introduced into the computing environment. These teams have operated with varying degrees of autonomy in managing their specific resources, and the process they developed and the tools they acquired have been specialized for those resources. These specialized processes and tools may not be appropriate when the organizational focus shifts from resource-oriented management to service-oriented management.

Because of the historical organizational structure within IT, moving from a resource orientation to a service orientation can present challenges in all aspects of the IT function. In general, the challenges of implementing ITSM span four broad areas: process, people, technology, and data.

Process

Implementing a successful ITSM strategy relies on the implementation of quality processes to provide IT services to customers. The correct development and implementation of these processes is critical to the overall success of the project.

For most IT organizations the implementation of true common cross-organizational management processes may be the most difficult aspect of the ITSM project.³ From an IT staff perspective, the organization probably does not change drastically with the advent of ITSM although new or altered responsibilities and organizational performance metrics are introduced. The introduction of new IT products and technologies and the phasing out of old products and technologies is a common occurrence; thus, the implementation of the required ITSM technical infrastructure is treated as “business as usual.” The implementation though, of day-to-day operational processes that span the organization may be a new experience for many of the IT staff.

Historically, IT organizations evolved based on the evolution of the technologies they supported. As new technologies were introduced, new organizations were put in place to support them. As the use of a specific technology increased within the business, the fraction of the IT organization supporting it grew accordingly. As these resource-

oriented organizations grew, they developed their own processes for supporting their resources.⁴

As common processes are developed and implemented across the IT organization, the following steps are required:

- Establishing process ownership
- Defining the scope of the process
- Agreeing on process design
- Developing process metrics
- Designing a technical infrastructure to support the process
- Deciding on process implementation priority
- Planning and executing process implementation and associated infrastructure

Best practices require that there be a single authoritative owner for each process. This person is responsible for identifying the process objectives, overseeing the process design, negotiating appropriate metrics, and eventual process implementation, operation, and improvement. Identifying the appropriate staff with the prerequisite knowledge, skills, and backgrounds to successfully own a cross-organizational process can be challenging.

Once a process owner has been designated, he or she should work out the scope of the process for which he or she is responsible. The scope includes all the resources, including hardware, software, and people, that will be under the control or influence of the process. Depending on the content of the services being provided and the process, the scope may also extend into the business units and the vendors who supply services to IT. Defining the correct scope of a process is important, and all affected groups should be involved in reviewing the process design and implementation plans.

There are differing opinions in the industry as to how to prioritize the implementation of the processes required for ITSM. They include the following:

- Start with the process or group of processes with largest business benefits.
- Begin with those that provide a rapid return on investment to the users of IT to garner continuing support.
- Implement configuration management and change management first, as all of the others depend on

configuration management, and configuration management and change management are closely related.

- Ensure that incident management is functional first, as this is the primary interface to the IT end-user community.
- Implement all processes in parallel, as the processes are truly interdependent.

Various consultants, analysts, and authors advocate any and all of the above. As a “best practice” has yet to emerge, the managers involved in a specific ITSM project have to decide on a priority scheme for their project. The prioritization should be reviewed with the executive management team for concurrence and then held to throughout the life of the project.

The successful implementation of a process in a production environment relies on fundamental project management skills. Detailed plans are required to educate participants in their roles, in the process flows, and in product usage. In addition, human and technical interfaces to previously implemented management products need to be identified and established.

The planning exercise becomes more complex when the process needs to be implemented across multiple sites. As with any significant project, poor planning that results in difficult implementation tarnishes the project before its initial operation. As the implementation of ITSM and its associated processes can conceivably affect the entire company, implementation planning is crucial to the acceptance of the new ITSM operational model.

People

One of the most difficult challenges in implementing an ITSM strategy is the impact ITSM may have on the IT organization and staff. The resulting changes range from shifts in organizational roles and responsibilities to changes in the day-to-day activities of individual IT staff. The identification, planning, and implementation of these people-oriented changes require time. Managers who implemented ITSM report that they spent 70–80 percent of their time working on organizational and staffing issues.⁵

Upon embarking on an ITSM strategy, understanding the true scope of the organizational changes required is imperative. Implementing an ITSM

strategy will obviously affect the direct IT staff, but there are other organizations and people that will be affected. This includes the business unit users of IT who now become customers, business unit staff who may currently provide some level of IT support, vendors who supply product support services, consultants who are supplemental staff, and business partners whose relationship with the company is heavily dependent on IT. The effect of a new ITSM strategy has to be evaluated and reviewed with each of these stakeholders before implementation.

For most IT organizations the scope of change that ITSM requires necessitates top-level management commitment. The great majority of IT organizations are hierarchical in nature. All reporting lines culminate at the chief information officer (CIO) or an equivalent role. Regardless of what might be touted as the strategy, hierarchical organizations encourage a behavior of “What’s important to your boss is what’s important to you.” Thus, the successful implementation of an ITSM strategy necessitates genuine, demonstrative commitment to the project from all levels of management. From the inception of the ITSM project, executive-level and middle-level management must display knowledge of and ongoing commitment to the project. Day-to-day project management may be delegated, but the IT staff must recognize that the commitment to ITSM is from the top down and that “it’s important to your boss.”

Once executive management is committed to an ITSM strategy, several organizational characteristics should be examined. Traditionally, the IT organization was shaped around the technologies supported; there was a network group, a workstation group, and so on. With ITSM the focus shifts to providing and supporting services, a new dimension. Decisions are to be taken whether to reorganize all or part of the staff to support services and to adjust the operating procedures and metrics of the current structure to emphasize the support of services. *Figure 1* illustrates a typical organizational change caused by the adoption of ITSM. Figure 1A shows a resource-oriented IT organization, whereas Figure 1B shows a service-oriented one. It should be noted that resource- or technology-oriented departments continue to exist, but there is equal emphasis on addressing the IT needs of specific business units through dedicated departments.




Figure 1
Examples of (A) a resource-oriented organization and (B) a service-oriented organization

Along with changes in the organizational structure, changes are required to the metrics used to evaluate the performance of the IT staff. In a resource-oriented organization, the staff evaluation criteria

may not directly relate to the ability of customers to receive IT service. For instance, no matter how good the mainframe availability is, if the network is not operational, the customer receives no IT service.

With the advent of an ITSM strategy, the organizational metrics should reflect the success of the organization in supplying services to its customers.

Many of the challenges in implementing a new strategy within an organization are due to the natural tendency of people to resist change.⁶ Carrying out organizational change in a high-pressure environment may not be welcome by the affected staff. For this reason any proposed changes should be thoroughly thought out, reviewed with all affected parties with all advantages and challenges identified, and carefully implemented in the least disruptive manner. If the IT staff understands the overall strategy as to why changes are needed and the significant advantages to implementing the changes, they are more likely to support the ITSM implementation strategy.

Technology

The technology required for an ITSM solution has three major components: (1) the configuration management database (CMDB), (2) process-level automation, and (3) task-level automation.⁷

CMDB

ITIL best practices strongly recommend the use of a CMDB to serve as an authoritative source of data.⁸ The CMDB is a repository of data concerning the entities to be managed, known as configuration items (CIs), and the relationships between them.⁹

Early implementors of ITSM generally embarked on developing their own CMDB, as there were no commercial CMDBs available on the market. Some were created based on an existing product database, such as a network management tool database; others were developed using standard database technology and products. Many of the early CMDBs that were tied to a specific product database turned out to be limited in their ability to support various types of resources as these were introduced into the ITSM solution. Today, almost every vendor of systems management products markets a CMDB product.

There are several aspects to consider when selecting a CMDB. This includes scalability in both scope and depth of information, data currency, and data integrity. Based on current technology, it is generally not practical for a large enterprise to rely on a single database that holds information on all CIs in the IT environment. In the extremely large, highly

complex IT environments of today, such a database could quickly hit scalability limits. Therefore, the ideal CMDB solution holds some of the CI data information physically resident (such as data associated with change management), and the rest of the data is held in remote repositories, where it is accessible on demand. This physical segmentation of data between locally stored data for quick access and remotely stored data should not be apparent to users. End users should be able to access all data through a single interface.

Process automation

There are a number of products on the market that can aid in process automation. The capabilities of these products range from products with process-specific support to generic workflow engines that can be customized. Many large enterprises have invested in what has become known as service-desk products. These products usually were conceived as automation engines for problem management or incident management, then were extended to change management and configuration management, and now some support a wide range of different processes.

With the advent of ITIL and best-practice process frameworks, products based on new technologies such as service-oriented architecture (SOA) have been introduced.¹⁰ These products are directed to the new ITIL-oriented market and advertise “out-of-the-box” (readymade) support of the ITIL processes and the ability to customize process flows as needed. An advantage of the SOA-based tools is their natural fit into a broader application environment.¹¹ In some cases the same application environment supports both business applications and management applications. This sharing of environments may have advantages in solution support and eventual integration of business and management applications over time.

As these products are new to the market, they should be examined carefully. Regardless of the product or technologies used, the ITSM process-automation infrastructure should contain some or all of the following characteristics:

1. Support for process customization at the task level
2. Support for defining roles and responsibilities
3. Interfaces to a CMDB or different CMDBs

4. Support for resource- or function-specific management tools
5. Scalability to the size of the environment, including the number of users
6. A rich set of status and performance metrics to enable process quality improvements
7. High availability and adaptability
8. Multiple levels of data and access security mechanisms

Task automation

As the new ITSM processes become well-defined, an evaluation of the installed toolset should be undertaken in order to determine current versus desired capability of process support. Depending on the current infrastructure, there may be tools available to aid in the execution of specific tasks within a process flow; for instance, software distribution tools to aid in change management, monitoring tools to aid in incident management, and performance-tracking tools to aid in capacity management or service-level management.

In some cases there may be several tools that perform similar functions, albeit the tools may be platform specific. For example, when teams of experts are dedicated entirely to a specific platform (such as UNIX**, Linux**, and Microsoft Windows**), it is most likely that each group uses platform-specific tools to perform operating system loads and software distributions. Once well-defined change-management and release-management processes are put in place, it is more efficient to have a single product that can serve all the platforms. This allows for easier integration (i.e., one instead of three integration points) as well as fewer external data sources that must be kept in synchronization with the CMDB.

Conflicts may arise should platform-specific support groups object to giving up platform-specific tools for a more generic cross-platform tool. However, the decision as to whether to keep a specific tool or to invest in a more general cross-platform tool should be governed by business needs.

Data

The implementation of a CMDB should take into account a number of factors.

1. Determining the granularity level of CI data
2. The reconciliation of information originating from different data sources

3. Determining what data should be stored locally and what data should be accessed remotely through database federation
4. Determining what data should be actively discovered and what data should be imported from existing data stores

Identifying CIs

The CMDB data model should be flexible enough to support the information required to describe all the CIs and their relationships. This includes both the types of CIs, such as “computer system” and “business application,” and the relationships between CIs, such as “installed on” or “depends on.” The granularity of CI data will vary based on business needs. ITIL suggests that it is generally best to describe CIs at the highest level that satisfies the business and service requirements. However, care should also be taken to understand the lowest level of CI that may eventually be required and ensure that the CMDB chosen is able to support this level.⁸ For example, an IT service may initially require only information such as “Database Server Y is installed on Computer System Z.” However, as process definitions continue to mature and understanding of the IT service increases, more detailed information is required, such as “Business Service W consists of Application X that depends on Database Server Y installed on Computer System Z.”

Reconciliation

CI-related data are often available in existing operational management data stores and may also be replicated in several different databases. The data for a specific CI may identify the CI in different ways. For example, an asset management tool may identify the computer system by its host name, whereas a network management product may identify the same physical asset by its Internet Protocol (IP) address. To provide a single authoritative source of CI data, the various CI-related data items must be reconciled and merged into a single standardized format. This is needed to populate the CMDB. The Distributed Management Task Force (DMTF) Common Information Model (CIM) sets the foundation for a standard modeling format.¹² Normalizing and reconciling the data into this format requires that any discovery or import mechanisms provided by the chosen CMDB allow for a rich set of rules that can be invoked as part of the import to ensure that only one unique instance of a CI exists in the CMDB and that the attributes

from the combined data sources are prioritized to reflect the most accurate data.

Federation

Once methods for normalizing and reconciling CMDB data have been determined, consideration must be given as to what data should physically reside in the CMDB and what can be federated from different data stores. Data federation is the ability to retrieve data from multiple sources and present the data as if they originate from a single data source.¹³

In general, CI attributes that change infrequently are candidates for importing and storing locally in the CMDB. For example, for a CI “computer system,” attributes such as “model number,” “serial number,” and “owner” should be stored locally. This type of information generally can be either discovered or populated from existing data sources and is an appropriate candidate for importing into the CMDB.

For data that are more volatile and change frequently, database federation is the more suitable approach. Database federation also is recommended for data whose currency and accuracy are important.

Situations with many users dictate that the data be stored locally in order to minimize federation overhead. Heavily controlled data, such as data under change management control, should be stored locally. Supplemental data, such as data beyond change management control, can be federated.

Populating the CMDB

Because populating the CMDB is one of the more difficult tasks involved in implementing ITSM, ITIL suggests that the initial load of the CMDB be automated as much as possible.⁸ Toward that end, there are three ways to collect CI data: (1) collection from existing sources, (2) auto-discovery, and (3) manual entry.

Most businesses have data stores that have been populated with systems management information by means of existing tools. In many cases, these data stores have accurate data about the environment and should be used to populate the CMDB. An example would be an asset or inventory database that holds desktop information. In this case, it may be desirable to automatically extract the data from the existing data store and populate the CMDB with this information. By using an open-architecture

CMDB, this data can be extracted and formatted in some type of standard format, for example, XML, and then loaded into the CMDB.

Extracting data from an external data source is also recommended if auto-discovery is not feasible. This may be a result of limitations in the discovery tool itself, because some physical devices are not reachable from the network or because of security constraints imposed on these devices.

For business applications and services, often the customer does not have a complete understanding of the components and relationships that make up the service and may require the use of auto-discovery tools. These tools can be either agent-based or agent-less. Agent-less technologies can be more accurate because agent-related malfunctions are not a problem. In any case, it is important that auto-discovery use open standards and require minimal security to perform its operations.

In general, there is no single approach for determining what data should be discovered and what data should be imported from existing data sources. Therefore, the practical approach to populating and maintaining the CMDB is a combination of auto-discovery, data import from existing reliable sources, and a minimal amount of manual entry.

APPROACHES TO IMPLEMENTING ITSM

Organizations that have begun to realize the value of an ITSM operating model have started down an implementation journey that is neither crisply defined nor clearly benchmarked in the industry. There are as many ways to implement a strategy as there are strategists and staff to implement their strategies. Each ITSM implementation has its own set of quick wins, challenges, unforeseen hurdles, benefits, and return on investment. Masamoto Yashiro, former CEO of Shinsei Bank, Japan, said of their transformation, “The real challenge of transformation was not in painting the end state but in choosing the means to reach it effectively.”¹⁴

To begin addressing the challenges described earlier in the paper, ITIL suggests working through four apparently simple questions.⁸

Where do we want to be?

In order to successfully implement ITSM, the business and the IT organization must be involved




Figure 2
Two approaches to implementing ITSM

in and committed to the transformation. The IT and the business units need to work together to agree on current requirements driving the ITSM project, outline the business objectives, and define a vision. Longterm benefits to all parties must be identified to encourage ongoing participation.

Where are we now?

To identify the changes needed in the IT organization and technical infrastructure in order to move toward the executive-established ITSM vision, the project team must understand their current state relative to the vision.¹⁵ A formal assessment or current state analysis should be conducted relative to the four aspects of ITSM: people, process, technology, and data.

ITSM self-assessment material is available from a number of sources. The IT Service Management Forum (itSMF),¹⁶ the organization supporting ITIL, provides free material on its Web site.

How do we get to where we want to be?

At this stage detailed planning to move from the current state toward the ITSM vision can be started. There are two basic approaches to plan and implement an ITSM strategy, top down and bottom up. These approaches are illustrated in *Figure 2*.

As can be surmised from the figure, the top-down approach is more process-oriented, driving toward the development of a technical infrastructure,

whereas the bottom-up approach may take advantage of currently installed products or products purchased with the intent to influence the development of processes. In reality, many ITSM projects develop a plan that has characteristics of both approaches.

Top-down approach

The top-down approach is generally advocated by consultants and service organizations. It is generally the preferred approach for large IT organizations and is characterized by the definition of required processes, followed by the design and creation of a technical infrastructure to support the processes. The approach may be preferable when

- formal IT management processes are already defined and operational at some level (there may be multiple versions of a process in operation),
- IT provides services to multiple customers,
- a greater return on investment is projected by standardizing on process, definitions, and organizational roles than on tool sets, and
- a common process design can be leveraged to define a common tool set.

This approach starts at the interface of the service and process layers illustrated in *Figure 2*. The IT assessment should identify what types of services the IT organization is providing to the business and what processes are used to provide those services. Required process improvements are then identified and prioritized to build the ITSM implementation plan.

Depending on the assessment methodology used, the plan may be developed through defining the current and future state of different characteristics and capabilities of a specific process. The matrix shown in *Figure 3* is a portion of a maturity matrix for change management. The rows are different characteristics of the process, and the columns are varying degrees of process maturity. In this example the current state assessment is plotted by using the grey diamonds outlined in red, whereas the desired state is shown through placement of the orange diamonds outlined in yellow. By prioritizing the processes that need improvement and then analyzing the desired improvements, a process implementation plan can be created.

As process designs are developed and approved by the IT organization and business units, suitable products and technologies are identified to support

	● Level 1 ● Unfocused	● Level 2 ● Aware	● Level 3 ● Capable	● Level 4 ● Mature	● Level 5 ● World Class
Accept and Classify Changes	Much confusion over the change entry process, or multiple (possibly changing) entry points. IT gets involved late in the cycle with no notion of authorization to request changes.	Clear entry points, but authorization process unclear and known to be frequently bypassed. Required information is not known by all.	Good enforcement of required information; tools/database used effectively; "informal" authorization process, possibly with some "rubber-stamping". Some RFCs are rejected early if data is missing or there is an obvious conflict of dates. Many RFCs are classified as "emergency" and allowed through the process.	Clear entry points; authorization works (evidence of some "rejects" or requests that need to be resubmitted due to insufficient information). Change manager confirms all priorities and categories. RFCs are always sent to the correct areas for assessment. Some emergency RFCs are due to poor planning.	Change entry is automated and process rules enforced as a result; lead times, process path, authorization requirements, etc. are always correct. Emergency RFCs are always justified and handled correctly.
Review Changes through Change Advisory Board (CAB) or Emergency Committee (EC)	Much confusion over the evaluation criteria: not agreed/communicated with the change process participants, and possibly changing depending on the assessors. There is no CAB.	Uneven enforcing of criteria and required information; process known to be frequently bypassed. Little or ineffective risk assessment. Lead times defined for major changes. Regular CAB meetings with a large group of people.	Evaluation regularly includes risk assessment. Lead times defined for all changes, but not enforced. Change types are defined but do not include all changes. CAB or EC sometimes limited to those affected by the change. RFCs sent out electronically for CAB preview.	Clear criteria; good balance struck between process standardization (automation) and meeting varied departmental needs. Risk assessments always done. Lead times required for all changes are enforced. Change types defined for all changes. Membership of CAB or EC always varies, depending on the RFCs being reviewed. Business areas may be represented on CAB.	High quality requests based on criteria that adapt to practical usage; organization feels positive about using the change management process; feedback loop in place. All change types accepted and controlled. CAB or EC frequently consider RFCs electronically without the need for physical meetings. Relevant business areas are always involved in CAB or EC decisions.
Authorize and Schedule Changes	Unauthorized RFCs are often implemented and there is no Forward Schedule of Change (i.e., change plan).	Some RFCs are implemented without being authorized and the change plan may be paper or a simple spreadsheet, or it may contain only major outage information.	Authorization and scheduling is done manually, though a diary system may be used to display the change plan. Schedule contains all major planned changes and most of the medium impact ones.	Some automation for authorization though change plan production is manual and can be 2 or 3 days behind reality. Change plan is online, is visible to most IT departments, and contains all changes.	High level of automation for authorization and scheduling. Change plan is available online for all to view and is always correct.
Current Desired					

Figure 3
Maturity matrix for change management

the processes. These products and technologies may be newly purchased or currently installed. In order to simplify the infrastructure, the services provided by the improved processes should be standardized. The use of multiple tools to perform the same tasks should be avoided.

Bottom-up approach

This approach is driven by the advantages in using current tools and their capabilities to help define the required processes. A quick return on investment is emphasized and achieved by minimizing the process design time and by quickly deploying a process-

supporting infrastructure. This approach is applicable in situations where

- the organization is relatively small and there may be no formal processes in operation,
- the number of customers or business units serviced is small,
- specific tasks are performed by a single tool or a small number of tools,
- the use of generic or minimally customized process designs is acceptable, and
- optimizing the investment in tools is important.

Consider for example a company in which the current asset management capabilities are poor. The asset management software used does not save details on data, relationships between data, or change history. It does not provide a business service view of assets, and its reporting capabilities are insufficient. To remedy these issues, the customer decides to implement a CMDB. It is decided that the IT processes can later be built on top of the CMDB.

As the implementation of the CMDB matures, there is the realization that a more formal change management process is required to ensure the integrity of the CMDB data. Not having any prior process definitions to work with, nor the need to invest in custom process design, the company purchases a change management product that comes with predefined process flows that can be customized to some degree.

A bottom-up approach includes the following challenges:

- The time involved in arriving at the point of providing ITSM standardized services may be significant.
- Tool and technology decisions are usually not of concern to executives, who may lose interest in the project.
- Project participants may get heavily involved in discussions of tool features and functions and lose sight of the long-term project goals.
- Elevating staff orientation to providing a service may prove to be more difficult as the approach is more concentrated on looking at the optimization of tool usage.

A bottom-up approach includes the following advantages:

- The project starts from a place where IT professionals are comfortable: tools and technologies and their use.
- The IT staff is more accepting of a direction that emphasizes improvement in current practices than the adoption of a new process.
- The new process minimizes costs because a base high-level design is used as a reference, and the details are developed based on what the tools can easily support.

The two approaches share some characteristics. They both require

- executive and upper management commitment,
- ability to provide tangible results, particularly in the short term,
- an ongoing, structured communication plan to continually advertise progress and generate enthusiasm for the project,
- a current state assessment and the development of an implementation plan to achieve the executive vision, and
- a dedicated project team to implement the plans.

The successful implementation of ITSM is an iterative process. As it is almost always a new operational model, designs and ideas need to be developed, tested, analyzed for improvements, and adjusted.

How do we know when we have arrived?

A key deliverable of the executive “setting the vision” activity is defining high-level project success criteria. These criteria may be focused on a variety of metrics or tangible accomplishments and include the following:

- Decreases in IT expenditures for operational support
- Creation of a supported IT service catalog
- Documented and operational IT processes
- A high percentage of IT services being supplied through supported service level agreements (SLAs)
- Tangible improvements in process metrics such as incident resolution times, change implementation time, and so on
- Implementing the CMDB and populating it with at least 90 percent of the required data.

Regardless of the criteria for success, it should have some direct or inferred business value. The criteria




Figure 4
Top-down approach for implementing ITSM at USD Financial Services

may change over the course of the project but any changes should reflect a continually supportive relationship between IT and the business units.¹⁷

IMPLEMENTING ITSM—AN EXAMPLE

We describe in this section a scenario that shows how a customer's implementation of ITSM might develop. Whereas the customer, USD Financial Services (USD, for short), is a fictitious business, the scenario is based on our experiences with a number of different customers. In this example we use the top-down approach to implementing ITSM and focus on the identification of requirements and the resulting design, from a top-level assessment to a product implementation.

Background

USD is an established multinational corporation that provides financial services for small and midsize businesses and personal finance markets. Different business units develop and support the following services: personal portfolio management, individual stock, mutual-fund, and bond trading, business loans, commercial real-estate loans, lines of credit, business brokerage, and currency exchange.

IT has been a part of USD since the early 1960s when it was used primarily for in-house administrative functions. Throughout the 80s and 90s, the advent of electronic financial transactions expanded the use of IT to business-to-business transactions.

USD has long prided itself on providing personal financial services to its customers. The preference among its client representatives is face-to-face meetings with clients. However, increasing numbers of customers are using Web-based applications and services to conduct business. These customers can conduct business any time. This increases the importance of software applications, which now must be always available and must have extended functionality to perform the sales and marketing functions previously performed by client representatives in face-to-face meetings. The change has also introduced a new role for the IT organization: supporting customers directly and not just internal users.

Recognizing the increasing dependency of the business on IT and the need to move to a service-oriented relationship with its internal and external users, the IT organization embarked on a strategic change to operating as a service provider. USD contracts with IBM to help in this transition.

Strategy

As USD has an established IT infrastructure, organization, and processes in place throughout its international operations, a top-down approach to ITSM was chosen. The corresponding major assets that IBM used in assisting this customer are illustrated in *Figure 4*.

The IBM team uses the Component Business Model* for the Business of IT (CBMBoIT) to analyze the

effectiveness of the IT organization in its various aspects. By using a business model for IT as a reference with the understanding that IT is expected to provide high-quality services to both internal and external customers, key components are identified for improvement.¹⁵

The IBM team identifies the configuration, change, and release processes of the IT operation as candidates for improvement. Once a plan is in place to address the immediate process needs, service definitions and workflows are discussed.

Defining management processes

Using the IBM Process Reference Model for IT (PRM-IT), a detailed analysis is conducted on how USD performs release management, change management, and configuration management. PRM-IT is an ITIL-aligned model that extends the scope of system management processes beyond the processes identified by ITIL. PRM-IT provides a high-level activity flow for each process and for the interconnections of all the processes in the model.¹⁵

Through the use of this reference model, the process activities that are performing well, those that require improvement, and those to be created are identified. Detailed process flows are constructed with the IBM Tivoli Unified Process (ITUP) Composer tool. This interactive tool provides detailed task descriptions for the processes described in PRM-IT, detailed role descriptions, process artifact descriptions, and tutorials on using specific IBM products to support processes.¹⁸

Once the various processes are developed, USD conducts evaluations in order to choose the products required to support the developed process flows: IBM Tivoli Change and Configuration Management Database (CCMDB) and IBM Tivoli Release Process Manager (ITRPM).¹⁹

Creating the CMDB

In the next phase of the project, the IBM team determines how the CMDB is to be populated with the appropriate CI data. In evaluating the current data sources of USD Financial Services, the IBM team notes the effectiveness of their asset management tool. The tool was developed in-house and had been in use for several years. It took the IT department considerable time to develop and train the business units in the use of this tool, which

became a part of everyday business activities. Therefore, the investment that USD made in this tool and the corresponding asset management data is to be leveraged in the configuration- and change-management processes.

The asset management process to be implemented collects basic hardware information for devices and other resources. It also collects resource attributes that are company specific, such as initial resource-cost and resource-depreciation schedules. It is determined that the IBM Tivoli Discovery Library Adapter (DLA) technology will be used as the interface for loading the data into the CMDB on a daily basis. This technology allows customers to leverage their existing investments in data.

To obtain a complete picture of the IT infrastructure in the CMDB, the dependency and relationship information between resources is available for visualization. To obtain this level of detail, the discovery function of the CCMDB is used. Using agent-less discovery mechanisms, the tool is able to retrieve the required CI dependency and relationship information. In addition, through the use of embedded reconciliation technology, the data entering the CMDB both from DLAs and through discovery are consolidated in a way that ensures that no duplicate CIs are created for the same resource. As a result, valuable data from USD's existing asset management tool was consolidated with business-system-application and CI-relationship information. Together, this information is stored in the CMDB, and can be viewed with the CCMDB topology graph. *Figure 5* shows an example of the topology graph for a specific business system.

The new CMDB also provides USD staff with the ability to perform change assessment and impact analysis during the access-change phase of the change process by visualizing the business-system and application dependencies associated with the change.

Automating process flows

Due to a need for rapid change-process and configuration-process performance, USD Financial Services requires an infrastructure upon which to automate the process flows created during the detailed process design phase. USD makes use of the Change Process Manager (PM) and the IBM Tivoli




Figure 5
Application dependencies graph

Configuration Process Manager (Configuration PM, for short) built into the IBM Tivoli CCMDDB and the IBM Tivoli Release Process Manager (Release PM, for short) to automate the process flows.

Each of the PMs allows USD to model processes according to the definitions that are documented in the ITUP Composer in process templates. These templates consist of a number of activities, each made up of a set of tasks that must be complete before moving to the next activity. Various task types allow USD to store process artifacts and information associated with a specific process.

The PMs also allow USD to create specific roles for users, such as change manager or release manager, and assign various activities and tasks to these roles. The roles are notified when a task is ready for action. The PMs monitor all of the activities and associated tasks and allow the company to understand and track exactly where the process is in its

life cycle and to understand if the process is waiting for a specific role or user to complete a task.

In the initial stages, USD decides to implement a pilot setup to model and test the process flows. The pilot is centered on a typical application upgrade scenario that was identified in the assessment phase and modeled in the ITUP Composer. This pilot goes through the complete change-process cycle, including invoking the release management process to build and deploy the release.

USD is quickly able to model the process based on the templates provided by the Change Management PM and the Release PM. By using the major-change-using-release template along with the standard-ITIL-release template, which closely match their process design, USD is able to clone these processes and then make the required modifications to model the application upgrade change-and-release-manage-




Figure 6
Change and release processes for the USD Financial Services application upgrade

ment process. *Figure 6* shows the activities that are modeled in the process flow. As can be noted from the arrows in the figure, when the Change PM gets to the use-release activity, the Release PM is automatically invoked and follows through until “Release is Distributed.” At this point, control is returned to the Change PM to review and close the RFC.

SUMMARY

In this paper we explored a customer environment that is representative of an IT organization striving to implement an ITSM strategy. We discussed some of the major reasons why businesses are now recognizing the need for ITSM, described some of the significant challenges that customers face when trying to implement ITSM, and discussed two different approaches to successful ITSM adoption. Finally, we provided an example showing how a customer might approach implementing ITSM by using IBM-specific experience and products.

ITSM as an overall IT strategy is gaining momentum in IT organizations that support many industries. This momentum is due to the service-oriented paradigm espoused by ITSM, which allows the business units and external customers to view IT as a service provider, much like other organizations they deal with and gives IT the responsibility to provide expertise in applying technology to business requirements. This clear delineation of responsibility provides an extremely clear and efficient operating relationship between customers and IT.

As companies adopt ITSM, improvements will continue in technology to support processes and service implementation. In addition, ITSM strategy adoption methodologies and techniques will continue to evolve. Companies such as IBM realize the immediate need for technology and methodology advances and are quickly stepping up to ensure their customer base is armed with best practice ap-

proaches as IT organizations increasingly adopt the ITSM vision.

*Trademark, service mark, or registered trademark of International Business Machines Corporation in the United States, other countries, or both.

**Trademark, service mark, or registered trademark of the United Kingdom Office of Government Commerce, The Open Group, Linus Torvalds, or Microsoft Corporation in the United States, other countries, or both.

CITED REFERENCES

1. A. Ganek and K. Klockner, "An Overview of IBM Service Management," *IBM Systems Journal*, **46**, No. 3, 375–385, (2007, this issue).
2. *Foundations of IT Service Management, Based on ITIL*, J. van Bon, M. Pieper, and A. van der Veen, Editors, Van Haren Publishing, Zaltbommel Netherlands (2005).
3. M. Jantti and K. Kinnunen, "Improving the Software Problem Management Process: A Case Study," *Proceedings of the 13th European Conference on Software Process Improvement, Joensuu, Finland, October 11–13, Lecture Notes in Computer Science* **4257**, Springer, Berlin (2006), pp. 40–49.
4. R. Schiesser, *IT Systems Management—Designing, Implementing, and Managing World-Class Infrastructures*, Upper Saddle River, NJ, Prentice-Hall (2002).
5. R. A. Steinberg, *Implementing ITIL—Adapting Your IT Organization to the Coming Revolution in IT Service Management*, Trafford Publishing, Victoria, BC, Canada (2005).
6. A. McAfee, A. Sjoman, and V. Dessain, *Zara: IT for Fast Fashion*, Product Number 9-604-081 Harvard Business School Publishing, Boston, MA (June 2004).
7. D. Lindquist, H. Madduri, C. J. Paul, and B. Rajaraman, "IBM Service Management Architecture," *IBM Systems Journal*, **46**, No. 3, 423–440, (2007, this issue).
8. M. Berkhout, R. Harrow, B. Johnson, S. Lacy, V. Lloyd, D. Page, M. Goethem, H. Bent, and G. Welter, *Best Practice Service Support*, The Stationery Office, Office of Government Commerce, London, UK.
9. H. Madduri, S. S. B. Shi, R. Baker, N. Ayachitula, L. Schwartz, M. Surendra, C. Corley, M. Benantar, and S. Patel, "A Configuration Management Database Architecture in Support of IBM Service Management," *IBM Systems Journal*, **46**, No. 3, 441–457 (2007, this issue).
10. C. Mayerl, T. Vogel, and S. Abeck, "SOA-Based Integration of IT Service Management Applications," *Proceedings of the IEEE International Conference on Web Services, 2005 (ICWS 2005)* New York, pp. 785–786.
11. C. Mayerl, F. Troscher, and S. Abeck, "Process-Oriented Integration of Applications for a Service-Oriented IT Management—Integrated IT Management Architecture," *Proceedings of the 1st IEEE/IFIP International Workshop on Business-Driven IT Management (BDIM 2006)*, New York, pp. 29–36.
12. Common Information Model (CIM) Standards, Distributed Management Task Force, Inc., <http://www.dmtf.org/standards/cim/>.
13. L. A. Haas, E. T. Lin, and M. A. Roth, "Data Integration through Database Federation," *IBM Systems Journal* **41**, No. 4, 578–596 (2002).
14. D. A. Upton and V. A. Fuller, *Information Technology and Innovation at Shinsei Bank*, Product Number 9-607-010, Harvard Business School Publishing, Boston, MA (November 2006).
15. M. Ernest and J. M. Nisavic, "Adding Value to the IT Organization with the Component Business Model," *IBM Systems Journal*, **46**, No. 3, 387–403 (2007, this issue).
16. itSMF International, <http://www.itsmf.org/>.
17. E. Šimková and J. Basl, "Business Value of IT," *Proceedings of the 14th International Conference on Systems Integration 2006*, Prague, The Czech Republic, <http://si.vse.cz/archiv/clanky/2006/simkova.pdf>.
18. C. Ward, V. Aggarwal, M. Buco, E. Olsson, and S. Weinberger, "Integrated Change and Configuration Management," *IBM Systems Journal*, **46**, No. 3, 459–478 (2007, this issue).
19. IBM Corporation, IBM Tivoli Release Process Manager, <http://www-306.ibm.com/software/tivoli/products/release-process-mgr/>.

Accepted for publication April 9, 2007.

Published online June 27, 2007.

Alan J. Keel

IBM Software Division, Tivoli, 3901 S. Miami Blvd, Durham, NC 27703-9135 (akeel@us.ibm.com). A senior consulting solution architect with the IBM Tivoli Advanced Technology Group, Alan Keel is a cross-brand solution architect with strong emphasis on IT service management. He is a Certified IT Service Manager with 19 years of experience in the area of systems and network management. Mr. Keel joined IBM in 1988, and over the past 11 years, has consulted for Fortune 500 enterprise customers on the design and deployment of systems management solutions involving both technical and IT-process implementation aspects. He has extensive expertise in the design and architecture of large-scale multinational systems management solutions, particularly with IBM Tivoli enterprise products and their integration with third-party vendors. During his career at IBM, Mr. Keel spent one year in Singapore (2001) and one year in Tokyo, Japan (2002), consulting with customers on the design and implementation of IT service management solutions.

Michael A. Orr

IBM Software Division, Tivoli, 3901 S. Miami Blvd, Durham, NC 27703-9135 (ormm@us.ibm.com). Drawing on his 27 years of experience in the IT industry, Michael Orr assists IBM customers and product developers in designing systems management solutions with processes defined to meet business needs and an appropriate technical infrastructure to support the processes. His unique background in product development and support and systems management process design enables him to create optimized customer solutions. As an early advocate of the Information Technology Infrastructure Library (ITIL) within IBM, Mr. Orr has spent the past seven years using this management process framework in customer services, product support, education, and most recently, in contributing to IBM ITIL-focused product development. Highlights of Mr. Orr's career include two hardware design patents, 100% Tivoli Sales Club for three consecutive years, and three levels of ITIL certification. He is also a certified instructor for the ITIL Foundation and Service Manager courses. His external activities include serving on the Cameron Executive Network at the University of North Carolina at Wilmington, and as a member of the university's Information Systems and Operations Management and Computer Science Advisory Board.

Robin R. Hernandez

IBM Software Division, Tivoli, 800 N. Magnolia Avenue Suite 1600 Orlando, FL 32803 (rohernan@us.ibm.com). The IT service management team lead for the IBM Tivoli Advanced Technology Group, Robin Hernandez has ten years of software development experience, with a concentration in data modeling, data warehousing, schema design, and database performance optimization. For the past four years as a lead solution architect in the Advanced Technology Group, she has worked with Fortune 500 enterprise customers to design, deploy, and customize IBM Tivoli systems management solutions. Her current technology focus is on data and process management integration requirements for successful CMDB deployment, and consulting for IBM software development teams on high-end customer requirements. Robin is a frequent speaker on ITIL-driven data and process integration at IBM customer and industry events.

Eduardo A. Patrocinio

IBM Software Division, Tivoli, 3901 S. Miami Blvd, Durham, NC 27703-9135 (eduardop@us.ibm.com). Eduardo Patrocinio is a senior IT architect in the IBM Tivoli Advanced Technology Group. He has over 15 years of software development experience, with expert-level credentials in the areas of database design and warehousing, network management, service desk tools, service level management and CMDB design. He is a principal consultant to internal IBM development groups on CMDB and ITIL process-management requirements. Over the last three years, he has implemented IT service management solutions for more than 50 customers in over 20 countries. His current focus is on the practical application of CMDB and configuration-management to IT service management and ITIL-based initiatives. His academic credentials include his selection as Engineering Student of the Year 1993 for the country of Brazil, based on country-wide testing and interviewing of students throughout the country.

John Bouchard

IBM Global Technology Services, 18000 W. Nine Mile Rd. Southfield, MI 48086 (john.bouchard@us.ibm.com). Mr. Bouchard, an IT Architect with IBM Global Services, develops IT service management solutions for Fortune 500 companies. Trained in IT service management at the manager level, he has spent the last 15 years working in various industries, including automotive, retail, telecommunications, insurance, and financial services, where he built solutions that leveraged IBM and other vendors' products. ■