Information Technology Customer Service: "Best Practices" Processes For Operations

M. Keith Wright, University of Houston – Downtown
Charles J. Capps III, Sam Houston State University

ABSTRACT

In today's information economy, information technology customer service managers require increasingly technologically advanced, but cost effective operations. These managers have a Byzantine array of standards, methodologies and best practice frameworks that offer hundreds of pages of detailed, but often confusing, guidance. This article attempts to offer a simpler guide for IT managers beginning a formal program to increase the effectiveness and efficiency of the day-to-day operation of their customer support services. The paper's goal is to summarize and clarify important key concepts from the literature, and also present lessons learned from consulting experiences with several large organizations. Finally, the paper offers ideas for future empirical confirmation of the recommendations summarized here.

INTRODUCTION

*D*oes your information technology (IT) organization have the capability to provide the required service with the required quality? If your organization is growing fast, there is a good chance it does not. When you as an executive finally make the decision to try to do something about these growing pains, the first step is usually management education. At that point managers encounter a bewildering array of standards, methodologies and best practice frameworks offering to help guide them. For example, Carnegie Mellon University's Software Engineering Institute publishes a process standard that can be used for many IT management processes. Called Capability Maturity Model Integration (CMMI), this new model is a process management standard, not a performance standard (Software Engineering Institute, 2002). Thus, the CMMI does not require organizations follow a particular process and is not meant to be an initial cookbook. Other standards focus less on the management process itself and more on the output of the process. An example is Motorola's Six Sigma programs which focus on defect reduction to a high level of statistical reliability (Motorola, 2007). In contrast, the COBIT (Control Objectives for Information Technology) approach focuses on auditable controls (COBIT, 2000). From a project management standpoint, there is the Project Management Body of Knowledge (Project Management Institute, 2004).

Yet none of the frameworks mentioned above gives detailed advice on the best practices for IT service management. The Information Technology Infrastructure Library (ITIL) was the first such framework (ITIL, 2000). However, ITIL provides detailed, but often confusing guidance for eleven processes for IT service management. ITIL and CMMI are distinctly different, but not mutually exclusive models. For example, the following CMMI practices correspond to some of the ITIL processes: Project Planning, Project Monitoring and Control, Integrated Project Management, Measurement and Analysis, Configuration Management, Requirement Management, Development Risk Management, Validation, and Product Integration. The related ITIL processes are: Problem Management, Incident Management, Change Management, Configuration Management, Service Level Management, Disaster Recovery, Security Management, Availability Management, and Capacity Management. The main difference between ITIL and CMMI is the latter's emphasis on software process maturity, whereas ITIL helps one understand and develop all of the areas of the IT infrastructure. ITIL has recently spawned a compatible international standard: ISO 20000 (ISO, 2005). This International Standards Organization (ISO) standard defines the requirements for a service provider to deliver managed services, and can be used as the basis for an independent assessment. It may also be used by businesses that are going out to tender for their services, to provide a consistent

approach by all service providers in a supply chain, to benchmark IT service management, or to demonstrate the ability to meet customer requirements.   This relationship between ITIL and ISO 20000 is often illustrated via a diagram like the one below:
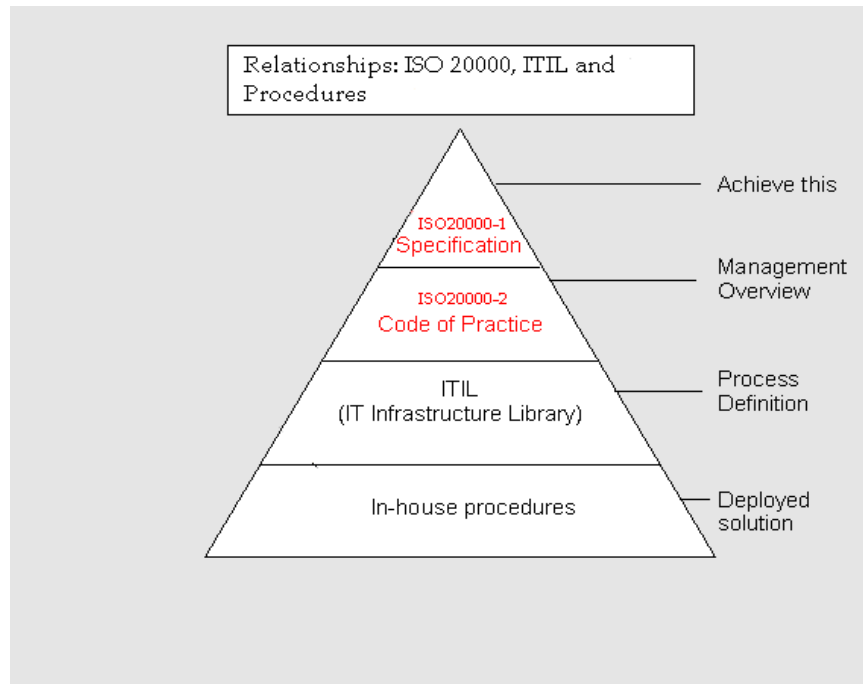
Today, Forrester (Garbani, 2006) estimates 30% of $1 billion-plus companies are experimenting with ITIL or ISO 20000 and between 12% and 13% have already implemented it.   Our consulting engagements revealed that many organizations have found parts of these models useful, but they struggled with problems caused by overlap, inconsistencies, and integration. Also, many organizations confront conflicting demands between these models and ISO 9001 audits or other process improvement programs.

The above literature encapsulates more than several hundred pages of tedious, detailed and overlapping guidance, gleaned from the best practices of large organizations over the past thirty years. The sheer volume of this information has created a demand for clarification in the form of training. To this end we provided IT consulting and training services that specialized in IT Service Management. Training courseware was developed to be compatible with the ITIL for two reasons. First, the ITIL provides the richest available source of information for IT service management.   Second, there are independent organizations that guarantee the quality of management training by means of independent testing and certification (http://www.exin-exams.com).

This paper is a synthesis of what was learned from conducting these training classes for hundreds of managers in various large organizations in the United States including IBM, the U.S. Department of Veterans Affairs, the City of San Jose, General Dynamics, Technology Concepts & Design, and Lowes Home Improvement Centers. In the training, we strictly adhered to ITIL, version 2 (v2), and ISO 20000 terminologies, because of the ITIL Practitioner's certification exam that EXIN administers on the last day of class. EXIN (Examination Institute for Information Science) is a global, independent ITIL examination provider. However, we found many of the EXIN ITIL v2 Practitioner exam questions, as well as parts of the ITIL v2/ISO 20000 literature are very confusing to many students. We believe this literature is needlessly complex. For instance, there are well known problems with ITIL v2 produced mainly because the ITIL book chapters were written by different authors, each imparting slightly different meanings to identical terms used in different chapters. We hope this paper improves the clarity of

ITIL v2/ISO 20000 by succinctly specifying its most important process areas and augmenting those with some opinions of our own. Furthermore, this paper, to our knowledge, emphasizes financial planning more than the published literature and suggests a practical sequence in which the process areas are best implemented a subject that ITIL v2/ISO 20000 does not cover. This approach means we occasionally use different terms than ITIL/ISO 20000 does, but these differences are explicitly emphasized throughout the paper. We caution that this paper is not suitable for managers studying for the ITIL Practitioner's exam because of our slightly different terminology. The key acronyms used throughout are listed in Table 1 below.

Table 1: Key Acronyms

SLMP	Service Level Management Process
SKMP	Service Knowledge Management Process
SDB	Service Data Base
IRP	Incident Resolution Process
ICP	Incident Correlation Process
SMP	Solution Management Process
CCP	Change Control Process
SLA	Service Level Agreement
SIO	Service Infrastructure Object

Our training course begins by framing ISO 20000 as an extension of ideas first popularized by W. Edwards Deming (Deming, 1966; Deming, 1986; and Deming, 2000). Deming observed quality business processes consist of five primary steps: plan, execute (do), measure (check), improve (act), and consolidate. (See Figure 1) The consolidation phase is when the organization ensures that improvements are documented and formalized through human resources, so they are repeatable.
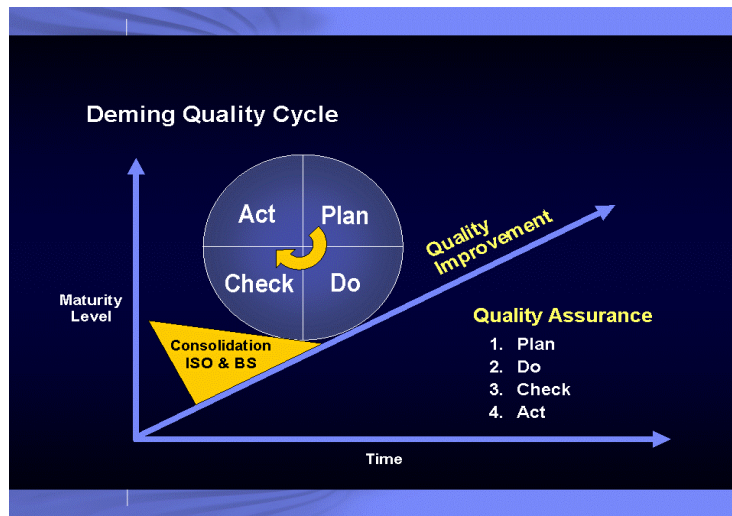

Figure 1

These ideas can be adapted to both small and large organizations for any business process. However, regardless of an organization's size, we believe the cornerstone of a successful quality improvement program is financial strategy.

Best Practice #1: Manage the Financial Expectations

A major challenge managers face when starting an IT customer service improvement initiative is maintaining positive expectations for the project over its entire lifecycle. Often these projects are very long term, and the benefits are not always obvious. The first step is to learn how much expense is justified. To do this, managers need to understand how much waste and/or lost opportunity exists in the current operation. This in turn requires someone be assigned responsibility for a separate and distinct IT cost accounting system. Because this is a full time job, this person ideally should not be the same person responsible for the organization's overall financial management. We call this new role the IT Financial Manager. This role should be charged with designing and documenting an IT Financial Management Process that provides accurate cost accounting to the other IT service managers. The inputs for this process are the business requirements, policies, and cost control methods. The IT Financial Manager develops a cost model used to allocate IT services costs to the customers receiving the services. Whether or not internal users are charged for IT services is optional although it is an excellent way to manage demand for IT services. Suggested cost elements are hardware, telecommunications infrastructure, software, payroll, transportation, relocation, overtime, consultancies, offices space, storage space, security, and utilities.

If the IT Financial Manager can prevent the actual costs of the initiative from exceeding the predicted costs, this will go a long way towards maintaining positive expectations. A steady stream of accurate cost-benefit analyses to stakeholders helps immensely. Earned value analysis is especially useful for this purpose (PMI, 2004). A word of caution: if the analysis does not show a potential for a huge benefit, go back to the drawing board. Quality initiatives are risky, even when kept simple. Once managers know how much money is to be spent on the customer service quality initiative, they are ready to get human resources involved.

Best Practice #2: Get Human Resources Involved

Before undertaking a process improvement initiative, we feel it is vital to get the human resources department (HR) involved, so compensation systems can be altered to include new responsibilities for essential personnel. ITIL/ISO 20000 does not cover human resource (HR) duties explicitly, but some HR roles can be implied from ITIL's mention of formal job roles including Process Managers and Process Owners. Combining this ITIL recommendation with ideas about COBIT's separation of duties controls, it is reasonable to recommend that these job roles be formally defined by an independent organizational unit (COBIT, 2000).

Best practices frameworks mention that a key prerequisite for process repeatability is Assigned Responsibility: there should be somebody accountable for defined processes covering objectives outlined in the financial planning effort. A formal project manager can make the design and implementation of any process more effective (PMI, 2004). The implementation manager may become the eventual "process owner." Once the implementation project ends and the process become operational, the "owner" remains responsible for the process outputs during its lifetime. From an audit standpoint, separation of duties is a consideration when making the decision to select the original project manager as the eventual process owner. (See Figure 2)

There is general agreement that processes, whether simple or complex, are always more effective if guided by the organization's strategy. Further, effective processes often cross organizational unit boundaries (ITIL, 2000).
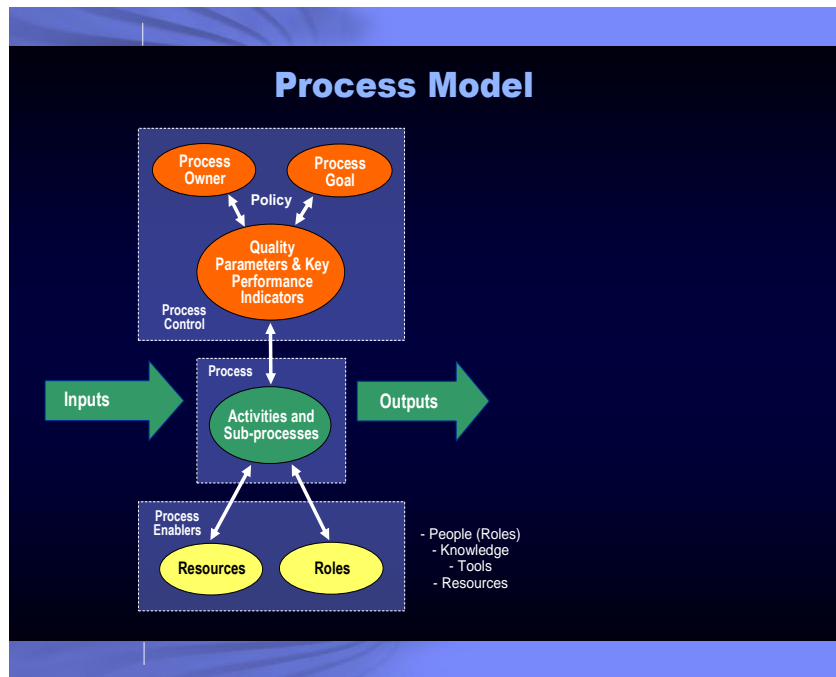
Figure 2

Best Practice #3: Create a Service Level Management Process

Once HR has established new job roles and incentive systems, we advise clients the next step is to implement what ITIL/ISO 20000 calls the Service Level Management Process (SLMP). Its purpose is to align and balance IT services with internal business needs, and to continuously improve the relationship between the IT services provider and its consumers. (See Figure 3)
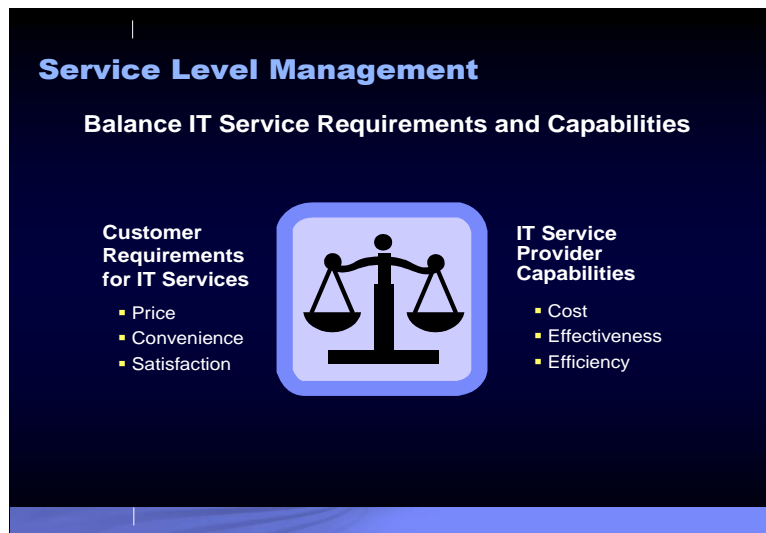

Figure 3

The "owner" of this process is often the same person as the IT financial manager. The SLMP is responsible for a written description of all services the IT organization is capable of and willing to provide for any consumer. Called the Service Catalog this document is analogous to the menu in a restaurant.

Inputs to the SLMP are the consumer requirements and the provider's capabilities. It is important to remember that consumer here means any user of the IT services, whether they are paying customers, or those who support paying customers. This includes users inside the provider organization. Their needs should never be ignored, because degradation in paying customer service may result (ITIL, 2000). According to many clients, IBM, once known for its royal treatment of employees, now has quite a different reputation; and struggles to keep its employees well supplied with information needed to support the paying customers. Often this results in frustrations and lost service business.

The principal output of the IT Service Management (ITSM) process is the Service Level Agreement (SLA), which is a clearly written statement describing the deliverables to a consumer including the agreed service hours, incident response times, service availability times, security requirements, continuity targets, and customer responsibilities (ISO, 2005). (See Figure 4) Although not a legally binding document, a clearly written SLA can avoid lawyers and courts by setting clear consumer expectations and promoting win-win trusting relationships with paying customers.
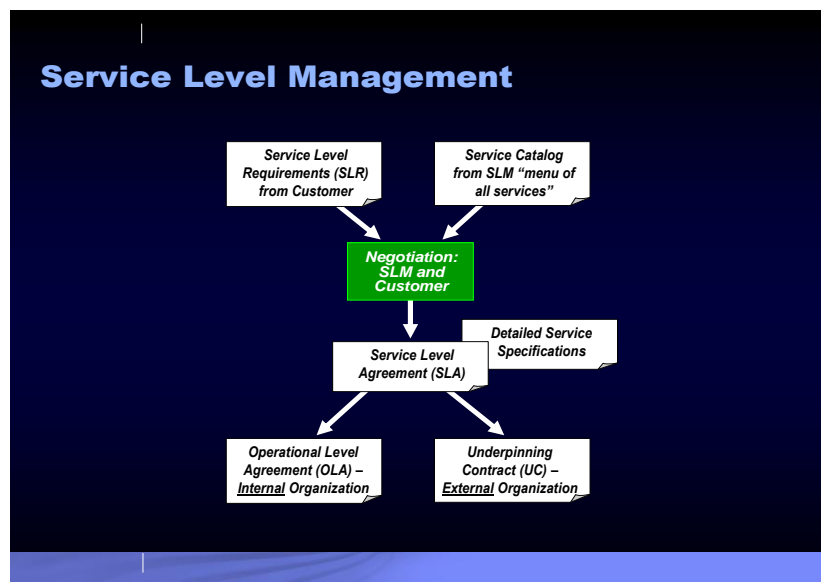

Figure 4

However, if the service level agreement depends upon external vendors, appropriately referenced underpinning contracts protect against legal liabilities (ITIL, 2000). Further, if the service level agreement depends upon multiple internal organizational units, appropriately referenced operational agreements improve communications among cooperating internal units (ITIL, 2000).

It is important that service level agreements have a section documenting the escalation procedures for use in case of service outages (ITIL, 2000). This section should distinguish between functional and hierarchical escalation, the former being cases where more expertise is needed, and the latter being when more institutional authority is needed. Functional escalation can be automatically triggered when agreed down-time intervals elapse. Hierarchical escalation procedures allowed any time during the incident resolution process can increase customer trust in cases where customer service cannot be restored within agreed lower limits. Automatic hierarchical

escalation can be considered when it is likely that timely resolution will fail. If service level managers can work out these details before the agreed resolution time is exceeded, line management can take early corrective action, for example by the hiring of third-party specialists.

Many clients report if service level targets are not measurable and monitored, stakeholders will likely not take them seriously, and service will not improve. Customer surveys are therefore recommended (ISO, 2000). Regularly held service level review meetings attended by paying customers, the service level manager, and other stakeholders provide an excellent forum for airing grievances and promoting a win-win relationship between the provider and consumer. This is the appropriate time for stakeholders to discuss possible amendments to the service level agreement (SLA).   Possible reasons for such amendments include additional charges for services not within the scope of the original SLA, such as office moves, major hardware and software roll-outs, unplanned hardware upgrades, etc. (ITIL, 2000).

Best Practice #4: Create a Service Knowledge Management Process

Once managers understand the costs involved and have assigned a single person responsible for service level management, the next—and most critical—process the IT organization must implement, in our view, is the Service Knowledge Management Process (SKMP). (We use the term service knowledge management in preference to the analogous ISO 20000 term, configuration management, which clients often find confusing.). This process provides complete, timely and accurate information to all the other stakeholders in the IT services organization. Accordingly, this process covers the identification, recording, relating and reporting of all service related infrastructure objects (SIOs) whether hardware, software, service level agreements, or other documentation. (We use the term SIO in preference to the ISO 20000 term, Configuration Item).

The inputs to the service knowledge management process (SKMP) are the service level agreements, other business objectives, and IT infrastructure details. The SKMP helps control the IT infrastructure through the identification, registration, monitoring, and management of all SIO documentation.

Thus, the principle output of the SKMP is the service database (SDB). (We use the term SDB in preference to the ISO 20000 term Configuration Management Database.) A well constructed logical model of the relationship among the SIOs, the SDB contains the kind of information shown in Table 2 below. Many SDBs can be initialized with data from existing asset management systems; this should be extended to include logical relationships between SIOs (ISO, 2005). For example, a disk drive is part of a PC. A PC is networked to a server. A program is a copy of the master. A manual describes a program.

Table 2: Contents of the Service Database (SDB)

Service catalog	Telecommunication services
Procedures	Service level agreements
Business units,	Servers
Supplier locations	Network components
Employee contact information	Desktops
Changes	Mobile units
Solutions	Software applications
Problems	Software licenses
Incidents	Facilities descriptions
Manuals	

An effective SDB enables support personnel to quickly identify all related service infrastructure objects pertaining to the SLAs. This in turn should improve service levels. Further improvements should occur if personnel can trace changes to service infrastructure objects (SIOs) from one state to another, for instance "under development," "being tested," "live," or "withdrawn."

A project manager plans the initial scope and level of detail to be included in the service database (SDB) (ISO, 2000). For example, s/he may not deem it economical to track each mouse and keyboard or to track SIOs proposed for future services. We recommend implementing the SDB incrementally, perhaps only including SIOs related to existing service level agreements. A full blown SDB software package can cost millions, so tools should be selected carefully. Inexpensive software can be used in a prototyping approach to help flesh out the details of the requirements for a major automation upgrade. Based on our experience, a cost effective service knowledge management process (SKMP) will eventually make or break the entire service quality improvement initiative.

After a simple SKMP is operational, its owner remains accountable for continuous accuracy of the service database (ISO, 2005). Thus s/he should control access to it and ensure no SIO record is added, modified, replaced, or removed without appropriate controlling documentation. Moreover, the SKMP should audit the SDB against the IT infrastructure after SDB implementation, after major infrastructure upgrades, and at regular intervals in between.

To briefly summarize, the outputs of the service knowledge management process (SKMP) are the service database (SDB) and the periodic audit results. (See Figure 5) The SKMP should start simply and become more complex as SLAs are added. Important SKMP performance indicators are the number of untracked SIOs detected in use, the number of SIOs used but not tracked in the SDB, and the number of changes to the SDB by type. Finally, if the SDB can save, protect, and report on historical time point snapshots quickly, it can be used to promptly restore the IT infrastructure to prior states, in case a change in the infrastructure causes unforeseen incidents.

Best Practice #5: Have a Service Desk Department

After the financial management, service level management, and knowledge management processes are defined and partly operational, we recommend the next step is build an appropriate service desk department complete with a budget for its staff. This department advocates purely for the consumer and is their single point of contact (ITIL, 2000). All consumer service requests, information, wishes, claims, complaints, requirements, quick fixes, and permanent solutions flow in and out of the service desk. The service desk often escalates incidents to other relevant departments and processes.

A customer facing staff skilled in communication and conflict resolution can improve customer relations and expectations enormously. Unfortunately, many individuals with career interests in computer technology are not suitable in roles that require customer interaction. Nevertheless, directly supporting the customer facing staff is a group of specialists with the required knowledge, skills, and authority to facilitate any kind of request or outage incident from a consumer. We believe these first level positions are excellent entry level slots for employees interested in careers in computing technology.

Without a service desk, consumers may not know where to request service and may interrupt technology specialists. Then outage incidents likely will not be reported completely and accurately. This prevents proper incident escalation and may cause the entire quality initiative to fail. Globally dispersed service desks can provide access 24 hours a day, 7 days a week. If so, it is vital that a common agreed upon language is used for incident recording, which is the first step in the incident resolution process.

Best Practice #6: Have an Incident Resolution Process (IRP)

Chiefly reactive, the mission of the Incident Resolution Process (IRP) is to restore service outages as quickly as possible, not necessarily to find the root cause of the outage. (We use the term incident resolution in preference to the ISO 20000 term, incident management.) Often this means simply reporting temporary quick fixes or workarounds found in the service database. An Incident Manager owns the IRP process. This manager is often the person responsible for the service desk budget (ITIL, 2000). An incident is an alert about a service disruption, or a potential service disruption (ITIL, 2000). Examples would be calls about printer failures from either a customer or internal support, or an alert about a disk overflow from a computer monitor. Efficient and effective incident reaction demands a formal method of working that can be supported by software tools.

First and foremost, the IRP accurately and completely records all incidents. Examples of data to record are incident ID, recorder, affected persons/systems, time, brief status description, and inventory number. The next activity is prioritization, which should be a function of three things: the category, urgency, and impact (ITIL, 2,000). Category and urgency are spelled out in the customer service level agreements. For example, category A48 might mean a service is now unavailable and must be restored within 48 hours to avoid payment delays. Category T3 might be a training incident that must be resolved within three months. Category S6 might be a request for an equipment relocation to be completed within six months. Additionally, the impact of the incident measures its potential effect on the entire service catalog. Although criteria for assigning impact is set up in advance via consultation with the provider's managers and formalized in service level agreements, the initial impact measurement is the subjective judgment of the service desk employees assigned to the incident.

Parallel with incident prioritization, the incident resolution process (IRP) addresses the incident queue in order of priority. At this point, initial investigation and diagnosis is performed. If an immediate solution is not apparent in the service database, the next step is to transmit escalation alerts to one of the target processes we detail later. Conversely, the IRP periodically monitors the output of the target process to determine incident status to report to the consumer. Examples of status include new, accepted, scheduled, assigned to specialist, impact, urgency, priority, work in process, on hold, time spent, specialists involved, costs involved, resolved, and closed. Support personnel should augment the incident description as the incident progresses towards resolution.

To review, the inputs to the incident resolution process (IRP) are incident reports from the service desk or monitors, the service database, service level agreements, and incident status. The outputs are escalation alerts, incident reports to the service database (SDB), routine requests for changes in the infrastructure, and status information to consumers (via the service desk). IRP personnel are responsible for closing an incident, but should never do so without the consumer's permission (ITIL, 2000).

Customer satisfaction surveys measure performance of the incident resolution process (IRP) (ISO, 2005). Possible survey questions include: Is the telephone answered courteously? Are services restored in accordance with the Service Level Agreements? Are you given good advice to prevent future incidents? Are you advised in a timely manner about current and future changes? Other performance indicators are the total number of incidents registered, the percentage of incidents resolved during the first alert, average incident resolution time, percentage of incidents handled within agreed response time, and number of incidents escalated correctly.

During normal IRP operation, incidents are escalated to the incident correlation process, detailed next.

Best Practice # 7: Have an Incident Correlation Process (ICP)

The primary mission of the Incident Correlation Process (ICP) is to constantly scan the service database vetting and grouping related incidents into problem records. (We use the term incident correlation in preference to the ITIL term, incident trend analysis, which clients sometimes find confusing). This can be automated via modern statistical and data mining tools. Employees in ICP roles are typically engineers.

Many clients come to us with informal incident management processes in place. These are often called 1st or 2nd line support. Most of these clients report their informal processes are ineffective.   This is due to two things. First, clients are often unaware of the latest data mining tools available. Second, they have trouble distinguishing between the following concepts: incident, problem, root causes, and defects. Part of this confusion arises from the terminology introduced by ITIL v2 and ISO 20000 where root causes are referred to as errors, and correlated incidents are called problems.

We advise that a problem is best conceived as the state of existing knowledge between the time point incidents are correlated (when problems are discovered) and a solution is discovered. Consider this example. Suppose a service desk employee receives a call about a printer failure. S/he would then record and prioritize an incident in the service database (SDB). Then suppose there are soon several more similar calls about different printers. An effective incident resolution process (IRP) proactively and constantly scans the SDB using data mining

tools and quickly discovers each of these failed printers are located on the same building floor. At that point a single problem record should be opened in the SDB. Subsequent similar incidents can then be added to that single record.

Therefore, the primary performance indicator for the Incident Correlation Process (ICP) is the ratio of incidents to problems. This ratio should be much greater than one. Without a formalized ICP the ratio of incidents to problems can approach one, and the sheer number of problems can overwhelm the entire IT service organization.

After the ICP associates incidents with problems in the SDB, the ICP prioritizes and escalates them to the next prescribed process, Solution Management. (We use the term solution management in preference to the ISO 20000 term, error control).

Best Practice # 8: Have a Solution Management Process (SMP)

The two purposes of the Solution Management Process (SMP) are to find the root cause of problems and then to recommend their solutions. Obviously personnel usually assigned to this process are engineers. To continue with our earlier example, suppose the engineers eventually discover faulty building wiring responsible for the printer failures on that building floor. The SMP then documents a defect in the SDB. (ITIL does not use the terms root cause, solution or defect, but rather the term error, which clients find confusing). The solution team then begins a complete analysis and documentation of a set of possible solutions.

The primary inputs to SMP are the escalation alerts from the Incident Correlation Process (ICP) and the infrastructure configuration details from the service database (SDB). The most important outputs are the solutions, workarounds, and request for changes (RFCs)—all recorded in the SDB. The RFC contains the set of possible solutions. An effective SMP also includes a proactive step to recommend RFCs for improvements to the IT infrastructure prior to any reported incidents.

The critical SMP performance indicators are the average time between problem entry and solution definition, and the ratio of successful to failed solutions. However, the eventual success or failure of solutions will not be discovered until the solution is implemented—by the Change Control Process which we discuss next.

Best Practice # 9: Have a strict Change Control Process (CCP)

ITIL v2 refers to solution implementations as changes to the information technology infrastructure. Changes to a high technology infrastructure can be extremely risky. However an appropriately strict Change Control Process (CCP) can effectively manage this risk. Thus, the CCP's purpose is simply to minimize the number of future incidents caused by changes. (We use the term change control in preference to the ISO 20000 term change management, which clients find vague.) The key to minimizing change related incidents is using a strict standard approval process for each solution, regardless of its nature. Accordingly, the CCP should take a strategic view of change to the IT infrastructure to ensure that all aspects of the change, both technical (deployment) and non-technical (training) are considered (ISO, 2000).

To continue the earlier example, an effective CCP ensures that change to the wiring affecting the fourth floor would not cause incidents on other floors. If IT organizations do not control change, they can be overwhelmed by it. Many of our Lowe's Home Improvement Center clients report that, in an aggressive effort to open one new store per month in the U.S., Lowe's in 2006 was overwhelmed by IT change related incidents, which often caused store IT systems to completely fail.

As a result, the first step in the CCP is to scrutinize and prioritize requests for changes (RFCs) input from the service database. RFCs should contain all information necessary for change managers to determine its priority (ITIL, 2000). At a minimum this includes an ID, change details, requester, change purpose, proposed priority, infrastructure objects affected, resource requirements, opportunity cost discussion, and related incidents/problems.

If the change request (RFC) does not meet minimum requirements, a change manager acting alone rejects the RFC for re-submittal to the solution management process. Further approval process steps depend on the category and priority of the RFC (ITIL, 2000). The approval process should always include a careful deliberation involving a committee of more than one person. (ITIL v2 calls this committee the change advisory board.) The size of the committee depends on the priority and impact of the change. At a minimum the committee consists of the change manager plus people involved in all the earlier described processes, especially the solution management engineers. Major change approvals always require senior executive sign-offs. Finally, the change control process (CCP) is responsible for planning, design, building, configuration, and testing of hardware and software prior to deployment. Testing should be done on non-production equipment (ITIL, 2000).

For these reasons, an effective CCP is often both capital and labor intensive. Many clients report that, in today's cost cutting economy, many (even large) organizations try to short circuit the change management process by testing on production equipment, perhaps during off hours. These short cuts often increase rather than reduce long term costs. Back-out plans—proven procedures to follow in case change related incidents overwhelm the service desk—further insure against failed changes (ITIL, 2000).

The change control process (CCP) schedules implementations and manages consumer expectations before deployment (ISO, 2005). Deployment of changes includes a sub-process called the Release Process (RP) (ISO, 2005). This important sub-process is responsible for grouping internally compatible RFCs into Releases, which are sets of RFCs that are best implemented at the same time. Complex software releases, for example, often involve intricate internal dependencies, best handled by specialists, and bundled into a testable unit. The RP should treat the release as a unit and record it as such in the service database. However the overall CCP approves the timing and contents of releases (ITIL, 2000).

The release process (RP) maintains and controls the Definitive Software Library (DSL)—a secure storage area containing all official, tested, licensed, and authorized software, along with associated documentation (ITIL, 2000). This storage area is kept separate and distinct from the development, test, and live storage areas. The DSL includes definitive copies of purchased software as well as internally developed software. The RP is thus responsible for all legal and contractual obligations for all hardware and software in use within the organization. The RP also maintains and controls access to the definitive hardware reserves. (ITIL calls this the definitive hardware store, which clients find confusing). This is an area set aside for the secure storage of authorized hardware spares. Its size depends on the scope of the quality initiative. These spare components should be maintained at the same version level as those in the production environment. The RP records details of the definitive hardware store (DHS) in the service database. These spares can then be used in a controlled manner when needed for additional capacity, or in recovery from major incidents. Once their (temporary) use has ended, the RP returns them to the DHS, or obtains replacements.

Finally, the change control manager chairs the Post Implementation Review (PIR) committee (ITIL, 2000). This committee meets periodically by teleconference to verify that previously implemented changes continue to resolve their related incidents. Members of this committee include representatives from all the process areas discussed, especially the engineers from the solution management process (SMP). Only after the PIR should the solution management process close the associated problem and defect records. The main outputs of the change control process (CCP) are successful releases to the production environment. Other outputs are the forward Schedule of Changes (FSC), Planned Service Availability (PSA), PIR meeting minutes, and action items (ITIL, 2000). The important CCP performance indicators are the number of change related incidents, number of releases deployed within budget, number of release failures, incidents of unauthorized software use, number of urgent changes, rate changes implemented, and the average cost per change (ITIL, 2000).

INTER-PROCESS COMMUNICATION

Now that we have briefly sketched an overview of the best operational processes in IT service management, it is important to note they execute in parallel with each other and thus require intricate inter-communication, depicted in Figure 5 below. Communication is both synchronous, via an alert system; and

asynchronous via the service database.   Many clients report difficulty in establishing effective communications between the incident management, solution management, and problem management processes regarding solution definition.   Smaller companies often attempt to deal with this by assigning employees to multiple processes. However, auditors may assess this practice as a separation of duties risk (COBIT, 2000).

IMPLEMENTING THE PROCESSES

We advise clients the best way to implement these processes is to take small steps in all processes concurrently, coordinated by one overall continuous service improvement program manager, with individual project managers responsible for their own process implementations (PMI, 2004). It must be stressed that implementing these processes can be very labor intensive in the short term. It is easy to underestimate the scale of work required. Effective financial and service level management processes are the foundations for all IT service management quality initiatives. Consequently, the service database (SDB) is of major importance.

Even if the organization is not purely an IT services organization, the change process must be managed from the highest levels of the organization (ITIL, 2000). Top management has the key role in avoiding the 'Silver Bullet Lifecycle' where new processes and increased accountability are introduced and then later abandoned during the critical initial period where costs are outstripping observable benefits.
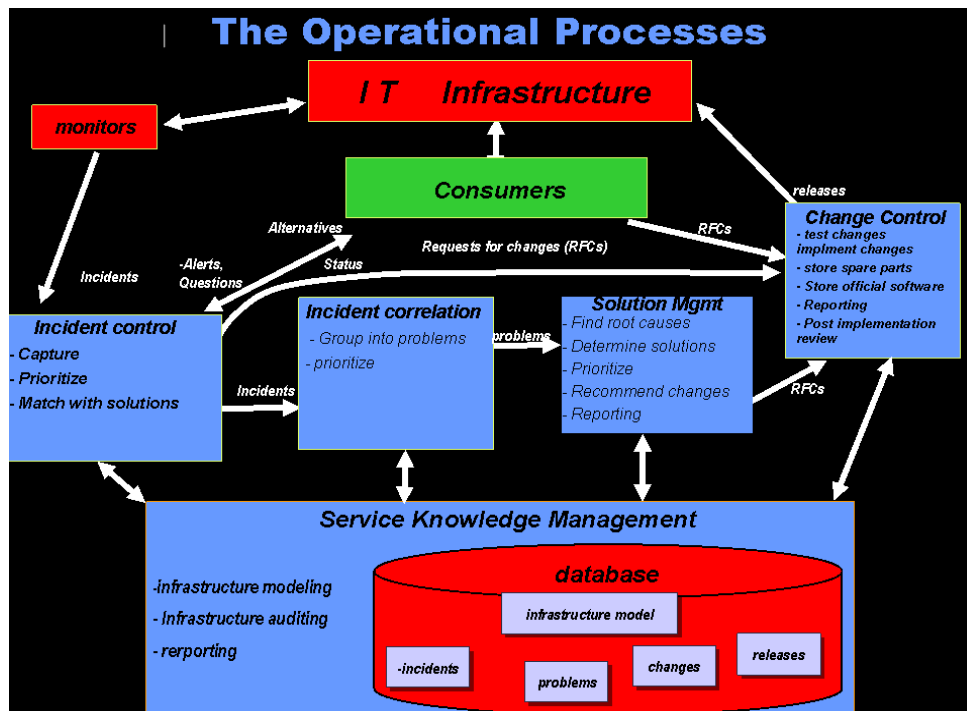

Figure 5

Regrettably, many clients report top managers do not like to get involved in IT service improvement programs because of the perception such programs are only the concern of IT specialists, rather than an integral responsibility of modern effective strategic management.

Despite the fact that poorly controlled infrastructure changes can fatally affect the quality of services, our recommendations to clients are often resisted. This is because excessive controls involve staff in unnecessary work. Thus we advise managers to base decision making on sound cost-benefit analysis, and to avoid overambitious

schedules. We find some of the resistance can be overcome with education. Furthermore, we advise that employees who follow best practice processes correctly should be compensated more than those that do not. In some cases, penalties may be needed as a last resort. For these reasons, we feel the human resource department must be involved from the beginning of the quality initiative.

Finally, we caution that automation tools cannot compensate for poor process management. Thus, we advise expensive automation tool purchases should be delayed until their requirements are fully identified. Often this is done via prototyping with either manual methods, or inexpensive existing tools. Many clients report managers and staff who blamed failed programs on tools instead of the processes.

CONCLUSION AND IMPLICATIONS FOR FUTURE RESEARCH

This paper offers a practitioner's overview of what the Information Technology (IT) community commonly considers "best practices" processes in customer service operations. These recommendations are grounded in the literature and our consulting experiences. This area of business research is becoming more important as IT services become commoditized like any public utility (Carr, 2003; Davenport, 2005; Iverson, 2006; Karmakar, 2004; Pal, 2005; Sambamurthy, 2005; Sawhney, 2004; and Sheth, 2006).

While this paper does not claim to be a comprehensive description of everything within IT service support, we believe it is a reasonable summary and extension of IT management's accepted "best practices." To our knowledge, many of these "best practice" processes have yet to be subjected to rigorous evaluation; consequently, we feel there is a need for empirical case studies that compare "IT best practices" companies with the rest of their industry. Empirical research is needed to further understand how these "best practices" benefit the bottom line. For example, do IT organizations who adopt generally accepted "best practices" see their IT services become more competitive? This is an expected result because IT service providers would have a single definable, repeatable, scalable, and consistent set of IT processes. Logically this should result in better organizational communications and in a long-term reduction in the cost of IT services, without sacrificing quality. Other questions that warrant further investigation include: do employees of "best practices" IT companies better understand their responsibilities and are they more motivated? Do customers of "best practice" companies see fewer service outages and faster response times? Are these customers more satisfied than customers of companies not following "best practices?" While we would expect this because service providers and consumers would better understand what is expected of them, empirical evidence would provide further insight. Perhaps the most important empirical question to pose is whether or not "best practices" companies see a reduction in internal business operations outages and a reduction in the number of urgent IT deployments of the kind Lowes Home Improvement Center experienced recently? Again, this is expected because hardware and software in use would be of known quality and more effective releases should enable businesses to cope with higher levels of change. Certainly these important questions need to be addressed by thorough future empirical research.

REFERENCES

1. Carr, Nicolas G. *IT Doesn't Matter.* Harvard Business Review, May 2003.
2. CMMI. Capability Maturity Model® Integration Version1.1 (CMMI-SE/SW/IPPD/SS, V1.1) Software Engineering Institute. Carnegie Mellon University. CMU/SEI-2002-TR-011, ESC-TR-2002-011, March 2002.
3. COBIT. Control Objectives for Information Technology, 3rd Ed. Information Technology Governance Institute. July 2000.
4. Davenport, Thomas H. *Coming Commoditization of Processes*. Harvard Business Review, June 2005.
5. Deming, W. Edwards. Out of the Crisis. MIT Press. ISBN 0-911379-01-0, 1986.
6. Deming, W. Edwards. Some Theory of Sampling. Dover Publication ISBN 0-486-64684-X, 1966.
7. Deming, W. Edwards. The New Economics for Industry, Government, Education – Second Edition. MIT Press. ISBN 0-262-54116-5, 2000.
8. Garbani, J. IT Infrastructure Library: Systems Management Essentials, December 4, 2006. www.forrester.com

9.   ISO. International Standards Organization. Service management-Part 2: Code of practice BS ISO/IEC 2000-2:2005.
10.   ITIL Version 2 (v2).   Information Technology Infrastructure Library.   Information Technology Infrastructure Library, Service Support. Office of Government Commerce, U.K. ISBN: 0113300158, 2000.
11.   Iverson, J. and Ngwenyama, O. *Problems in measuring effectiveness in software process improvement: A longitudinal study of organizational change at Danske Data*, International Journal of Information Management 26, 2006.
12.   Karmakar, U. *Will you survive the services revolution?* Harvard Business Review. Volume 82 Number 6: 2004.
13.   Motorola Corporation. What is Six Sigma? www.motorola.com/content.jsp?globalObjectId=3088, 2007.
14.   Pal, N., R. Zimmerie. *Service innovation: A framework for success. eBRC* White Paper, Penn State eBusiness Research Center, SMEAL College of Business, Pennsylvania State University, University Park, PA., 2005.
15.   PMI. Project Management Institute. A Guide to the Project Management Body of Knowledge-Third Edition, ISBN: 193069945XISBN13: 9781930699458, 2004.
16.   Sambamurthy, V. Editorial Notes. Information Systems Research. Volume 16, Number 2, June 2005.
17.   Sawhney, M. and Balasubramanian, S. *Creating growth with services,* Sloan Management Review. Volume 25, Number 2, 2004.
18.   Sheth, A., Verma, K. and A. Gomadam. *Semantics to energize the full services spectrum,* Communications of the ACM, Volume 49, Number 7, 2006.

NOTES