IT Infrastructure Library (ITIL ®)

Sue Conger, University of Dallas, USA

IT Infrastructure Library (ITIL ®)

ABSTRACT

The Information Technology Infrastructure Library (ITIL ®[1]) is a series of processes that are required to run a quality IT operation that delivers value to its parent organization. ITIL has become important to businesses that seek to align their IT operations with the business' strategy. This chapter outlines ITIL version 3, which was published in May, 2007 and discusses issues in its application to real-world situations.

INTRODUCTION

Businesses have long sought to control their IT applications development practices with the implicit assumption that such control would lead to maintainable and operable applications that produce value for the business. However, in the 1980s the UK Government embarked on an exercise to improve its operational functioning in the hope that it would improve the value of IT to the government. The outcome of the UK exercise was the first version of the IT Infrastructure Library (ITIL ®), a series of books that document best practices in the management of the IT operations function.

Process management is important for many reasons but one of the most compelling is that it leads to greater payback for an organization. As one 2005 study shows, organizations that manage neither IT nor process gain no payback from either

[1] ITIL is a registered mark of The Office of General Commerce of the Government of the UK. It is usually written as ITIL® but the ® registered mark is omitted hereafter for readability. If this violates the rules of registered mark usage, it will need to be added back.

(Dorgan & Dowdy, 2005a). Organization that use IT to support their business processes but that do not explicitly manage their processes can expect a 2% return from the IT investment. Firms that manage processes with low IT support can expect an much as an 8% return while those that manage both process and IT investment support can expect a 20% return (Dorgan & Dowdy, 2005b).

Managing the IT investment is more than just managing development; and, as companies buy commercial off-the-shelf software, development plays a lesser role in the organization's total cost of ownership for IT. IT organizations minimally are comprised of several organizations (see Figure 1), including application development, database

Figure 1. Simple IT Organization Structure

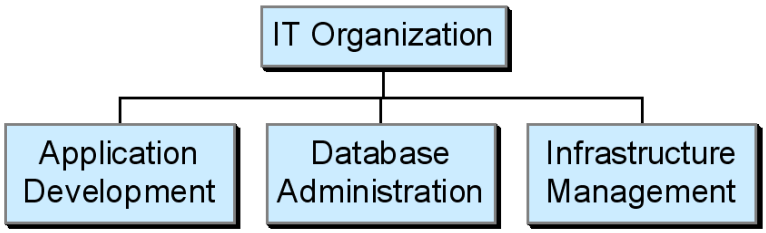

administration and infrastructure management.

Infrastructure management includes not just management of the physical computing equipment such as routers, servers, tape drives and the like, but also the efficient and effective operation of the infrastructure. ITIL addresses the infrastructure operation management task.

IT operations is critical to organizational effectiveness since as much as 90% of IT budgets now are used to manage operations (Fleming, 2005). **Service management** is generally used to refer to the management of processes within IT Operations so that, through efficient and effective execution of the processes, value accrues for the parent organization. Thus, companies are now recognizing that value can be created through application of best practices to IT Operations.

The term "**service**" has no single definition and ranges from a change in condition or state of an entity caused by another to a set of deeds, processes, and resulting performances (Zeithaml & Bitner, 1996). From one perspective, a service is "work done by one person or group that benefits another" (Farlex, 2008). From another perspective, a service is the execution of a process (Fitzsimmons & Fitzsimmons, 2005). From the 'official' ITIL v3 perspective, a service is "a means of delivering value to customers by facilitating outcomes customers want to achieve without the ownership of specific costs and risks" (TSO, 2007, p. 45; TSO, 2007).

While service management begins with IT strategy, the heart of service management is a series of 10 processes and one function (service/help desk). "A **process** is the set of activities (repeated steps or tasks) that accomplishes some business function" (Conger, 2008, p. 4).

The processes in ITIL all relate to keeping an operations organization functional. The ten ITIL processes relate to management of incidents, problems, change, release, configuration, availability, capacity planning, financial planning, continuity, and service levels. Hence, ITIL tends to be implemented in the infrastructure organization first. However, many ITIL processes, for instance, incident and change management, are initiated within operations but are actually remedied or executed within another organization usually within IT. Thus, service management processes have tendrils that permeate other IT organizational processes. Therefore, coordination of activities throughout an IT organization is needed to ensure successful and encompassing ITIL implementation.

While ITIL is the only best practice framework that principally addresses IT Operations, there are many valuable alternatives to ITIL that a company might adopt. For instance, the Control Objectives for Information and Technology (CobiT©), the framework most closely related financial reporting compliance (e.g., Sarbanes-Oxley in the U.S.), initiated in the auditing world but has crossed over to management of the IT organization. Another often-used framework is the Capability Maturity Model – Integrated (CMMI©), was originally developed to support application development management but has crossed over to use by operations organizations for such areas as project management.

Similarly, there are customized versions of ITIL by Microsoft – the Microsoft Operations Framework (MOF©), Hewlett-Packard, IBM, and others. These frameworks adopt ITIL as their base and build on them by customizing for a suite of support software

that imbeds the process within the operational framework of software for help desk, network monitoring and the like.

For all of these frameworks, and for service management in general, the goal is creation of value to the organization through its IT operations function. A secondary goal and outcome of successful service management implementation is alignment with the strategy of the business.

While all iterations of ITIL have had alignment as their goal, there has been clearer progress toward articulation of exactly how alignment might be achieved with each generation of the framework. ITIL version one (v1) was a loose collection of 30+ books each covering a different problem area within IT, spanning development, management, and operations. The scattershot approach was more focused in version two (v2) which reduced the number of books to seven, integrating much of the material from the 22 books. V2 of ITIL also focused more specifically on IT operations with the 'red' and 'blue' books on service support and delivery that articulated the ten processes and their best practices (Cartlidge et al., 2007).

Version 3 of ITIL added two key topics and rearranged material from version 2. The new topics are ITSM strategy and the life cycle concept as applied to ITIL practice. Details of these are provided below, but the main outcome from their inclusion has been to provide a clear starting point for evaluating changes to ITSM practice whenever a strategy is created or modified. The life cycle provides another avenue for re-evaluation of ITSM practice when, during continuous improvement exercises, a process is evaluated for change.

Some items noticeably missing from all versions of ITIL are generic process maps of each process, specific, step-by-step guidance on key steps in each process, examples of what some companies did, a differentiating of required versus optional components to say one has accomplished process management in an area. These shortcomings are compensated for by consulting companies which, for a fee, do have generic process maps, and by the organization IT Service Management Forum (ITSMF), a global non-profit group of practitioners who advocate ITIL practice.

This chapter is structured to summarize each of the major topic areas within ITIL v3. First, the service life cycle is articulated to describe its purpose and practice. Then, service strategy is more fully articulated. Service design follows and describes both the infrastructure and architecture decisions made to build-out an IT operation and also the five processes that are needed to support maintenance of continuous operations. The service transition section describes the activities needed to introduce changes into the IT infrastructure, ranging from management of incidents that remedy an outage through to configuration management and the relationships between configuration items. A summary of roles and responsibilities is provided. Finally, the material is summarized briefly with a discussion of ITIL 's role in actualizing 'service management' in IT organizations.

THE SERVICE LIFE CYCLE

Version 3 of ITIL introduces the service life cycle. The cycle formalizes the plan-do-check-act activities discussed in version 2 but not explicitly part of ITIL ®. The life

cycle begins either with the development of a new or changed business strategy or with a continuous improvement review that determines that changes are warranted.

Within the life cycle, policies and processes needed are the first decision. For the processes, the level of service required is defined. Then, the new or changed service requirements proceed through a design process to define sourcing, risk management approach, interfaces, and metrics. Once designed, the new or changed service is developed and quality assured through the transition processes, eventually being placed into operational status. Once operational, the new or changed service is monitored and reported on to ensure that is meets its utility and quality goals. From the metrics, either the service remains in its present status, or it undergoes further refinement and another iteration of the life cycle stages.

Policy & Process

A **policy** is a statement of intent by the executives of an organization. A good policy contains elements of purpose, reason, goals, and compliance. Policies are defined during the IT Service Strategy phase. A policy says 'what' will be done; a process defines the 'how'. A **process** takes inputs from vendors and customers and uses IT resources to conduct the steps of a computing service that is measured. Outputs flow to the customers either real-time in the form of computing resources or post-hoc in the form of reports based on the computing process. Based on the measures, adjustments to improve the process are made. Processes are defined during the Service Design phase.

Service Level Definition

A service level definition defines the service value in terms of the outcomes expected (the *utility*) and the level of IT support in terms of response time, capacities,

security, risk mitigation, continuity, hours of availability, and so on. The definition also describes the provider type as dedicated, shared, or utility, and internally provided or outsourced. Critical success factors identify areas to be measured to ensure that the service level achieves the desired outcomes. If interfaces with the customer or other organizations are required as part of the service level definition, they and their level of service support for all of the above are also defined. The outcome of service level definition is both a blueprint for the design of computing resources and services, and also a contract between the IT and client organizations.

Service level definitions are defined partially during the Service Strategy phase and partially during the Service Design phase. During strategy formulation, the business requirements in terms of response time, hours of operation, accuracy, reliability, security, privacy, and so on are all defined in business terms. During Service Design, those business requirements are translated into technology resources and costed to provide estimates of incremental and ongoing cost for meeting the business requirements.

Also in the design phase, once the definition and costing of all capabilities and resources is complete, the service level definition becomes a **service level agreement,** or contract, between the IT organization and the customer. Details of this process are enhanced in the following sections.

Design Process

The design process is the phase during which existing capabilities are mapped to the new or changed service level definition to determine whether or not new capabilities are needed. The technology, application, and data architectures for the new or changed service are integrated with organizational architectures to determine 'fit' and the need for

cross-organizational negotiation to ensure architectural synchronization. Once the architecture and technologies are resolved, the design progresses to create secure and robust operating environment with controls and metrics. This phase is more fully defined below.

Transition

Once design is complete and IT resources are identified, transition processes supports management of the changes to accomplish acquisition and IT resource allocation to the new or changed service. When the changes, either to applications, processes, or technology platforms are complete, other processes are initiated.

The release process ensures adequate testing, an operational back-out capability, and readiness for production. Working with the Change Advisory Board and Change Manager, the Release Manager certifies a change and schedules its transition to operational status.

With these basic definitions and overall understanding of the service life cycle, we turn to defining the individual phases of the life cycle beginning with service strategy.

SERVICE STRATEGY

During the development of an ITIL service strategy, the goal is to define services that support the parent organization's key initiatives, providing sufficient value through the service's utility and warranty. Utility defines the desired functionality and resources required while warranty defines the delivery strategy in terms of availability, security, and continued operation (Cartlidge et al., 2007). Critical success factors are defined for each service unit to monitor the service's operational delivery.

Before developing a service strategy, the target for service definition is assessed to determine its maturity and readiness to accept an IT service. Context, in terms of both industry and company size, are considered in the definition along with organizational goals that define the competitive stance of the organization desired.

In developing the IT strategy, there is a necessary tension between recommended best practice and business realities that need to be balanced. Business realities include such items as resource constraints, market position of the desired service objective, and service alternatives and costs. In addition to business context, a risk profile for the service is developed to define the extent to which high-availability computing, security, and compliance monitoring and tracking are required.

The outcome of strategy definition is a service-provisioning model that defines the type of service delivery as a sole-service managed activity, shared service, or utility service. Shared services include infrastructure and IT resources. Utility services are provision of on-demand service such as eMail.

Strategy Definition

To begin defining strategy, one must first understand the context and customer for which the strategy is being developed. Adopting Mintzberg's 4Ps of strategy, one first understands the customer and product being supported from four perspectives: Perspective, position, plan, and pattern (Nieves & Iqbal, 2007, pp. 91-93). By defining perspective, the "vision and direction" determine the overall positioning of the product in the marketplace (Nieves & Iqbal, 2007, p 92), thus defining the desired state. Then, the marketing 'position,' such as 'low cost' is defined. The plan articulates the tasks needed to move from the current state to the desired state. Finally, the customers needs for IT

resources is expressed in 'patterns' for utility, shared, or custom services, for instance, customized high-availability computing. Customer outcomes are often expressed in terms of improving capabilities, resources, or performance, or in terms of decreasing costs or risks. These definitions guide the service definition and also the metrics needed to measure their success.

An example of an IT service strategy statement might be "X desires to improve computing availability with outages no longer than 15 minutes for the ERP system used company-wide and no security breaches from outside the company." There are several goals here: Improve availability, manage outages, and no security breaches. Each of these would be analyzed in terms of the cost and approval of expected expenditures approved in principal. At this point, the costs would be only estimates.

By decomposing the overall strategy goals, each is modularized and uncoupled, allowing independent analysis and implementation planning while still needing overall integration to ensure that the overall goals are accomplished.

Each goal is analyzed a bit further to create a 'system' of inputs that feed process steps that create and deliver outputs that generate customer feedback. Any negative feedback is to be used to improve either the quality of the inputs or the efficiency and effectiveness of the process steps.

Once goals are fully understood, the IT services to support the goals are initially identified and a financial analysis is developed to estimate the budgets for development, transition, and for on-going operations.

A final aspect of strategy is a high-level of demand management. Strategy flows from the executive board of the organization and often the board requests multiple goals

for which prioritizing of IT resources is needed. The executives work with the CIO to understand the expected IT resources and the associated costs for each requested goal. Then, when there are scarce resources, the executive board prioritized the service offerings for development and delivery.

Financial Management Design. For any strategy, a business case and financial payback estimate from the strategy should be developed. Issues such as market differentiation, costs, customer outcomes, priorities, efficiencies and inefficiencies are related to services with costs and benefits identified. Service pricing as "cost-to-value" provides a cost baseline for services to provide IT operations visibility and improve consumption behaviors (Nieves & Iqbal, 2007).

To arrive as a financial valuation, the two components of service value – utility and warranty – are decomposed to develop a financial cost. Inputs are actual costs of licenses, maintenance, equipment, personnel, overhead for plant and equipment, taxes, compliance costs, write-offs, depreciation expenses, and so on. The value obtained is expressed as the monetary value of individual service components, e.g., service desk support based on number of calls for support, number and management of outages, etc. relative to the total for the service desk. Monetized elements include costs for staff, hardware, software, facilities, and capital (Nieves & Iqbal, 2007). This section of ITIL is not well articulated in that it offers little guidance on how to actually develop the values for services. Cases suggest benchmarking costs against those of other companies, e.g., for the service desk, and ensuring that all costs both for acquisition and disposition of assets be included in financial analyses.

SERVICE DESIGN

Once a strategy is defined, service design begins. "Service design starts with a set of business requirements and ends with the development of a service solution" (Cartlidge et al., 2007, p. 18) and is comprised of three steps. First, the organization evaluates the new or changed service relative to its existing service portfolio and a revised (or new) set of architectures for hardware, software, and data are developed as required. Then, the hardware and software components of the service are designed. Next, the tactical processes for operating the environment are evaluated for change impact. From these evaluations, each with its own output, changes to existing services and processes are considered and modified as needed. Finally, metrics are designed to prove that the service is meeting its utility (i.e., functionality) and warranty (e.g., availability, security) requirements. Each of these steps is discussed in more detail. The outputs of the service design phase are architectures, policies, processes and documentation to meet current needs while retaining sufficient flexibility to scale into the future (Nieves & Iqbal, 2007).

Service design results in several different outcomes: Architectures, definition of the service portfolio and catalog, design of the new or changes services and all components thereof, concomitant design of processes required to support service provision, and design of metrics to prove service success.

Architecture Definition

In ITIL ®, an **architecture** is the structure of components, relationships, and underlying principles to accomplish a function or functions (Cartlidge et al., 2007; Nieves & Iqbal, 2007). Architectures can be technical, such as the organization of inter-relationships of modules within a hardware system. Architectures can be more

encompassing, including roles and responsibilities from business, vendor, user, and IT communities, plus the interplay of hardware, software, and databases, plus the IT service processes needed to support the business, plus metrics and a detailed definition of service levels expected. Both definitions might apply in the ITIL context.

Some underlying principles might be "using software as a service (SAAS)," or "supported by SAP ERP," or "applying current organizational databases as appropriate." These help guide the architect in evaluating the appropriate 'current state' related architecture and determining the extent to which changes are required to accommodate the new or changed service.

Just as strategy considers the 4 Ps of Strategy, service design considered the 4 Ps of Design: People, product, process, and partner. Each of these is considered in the definition of the architecture for the service(s) to be provided. In addition to the architecture for the target service, individual architectures within the organization for data, operating environment, software and processes may be analyzed and either developed or updated. These architectural exercises are related to organizational maturity in its management stance toward IT. The more mature the organization, the more likely they are to have architectural management as part of their IT service structure.

Architectures can be simple, such as an entity-relationship diagram as a data architecture. Or, architectures can be complex, multi-component depictions of roles, processes, and technology.

Tactical Process Design

Service design includes definition of IT technology components, processes, and metrics. IT components can be alternatively thought of as technology assemblies that are

imbedded within a process, an organizational capability or resource, categories of services each with a related cost, or asset streams that contribute to the revenue-production of the firm. Each of these perspectives is evaluated during service component design

Components. While there are several different ways to think of components in ITIL ®, the most common is as hardware/software assets; this perspective is developed here. The proposed service is examined to determine the extent to which functional needs are met by existing software, data needs are met by existing databases, and hardware needs are met by existing infrastructure. In mature, large organizations, this evaluation might be accomplished by means of statistical modeling for infrastructure alternatives. In small or immature organizations, this might be a paper/pencil exercise in which documentation is evaluated.

If new software is needed, a request for change is initiated and requirements for the software component are developed. If infrastructure – hardware, telecomm or other computing support – is needed, a capital expenditure request is initiated and the business case for the expenditure is developed with the business unit client. If new data are needed, a request for change of the database is initiated and the database administrators work with the application developers (or software evaluators if purchasing software) to define conditions under which data will enter the database, validation, security and privacy requirements, and any other controls needed.

The output of the component design part of service design are the functional requirements for data, application, or infrastructure changes and the initiation of those activities.

Process. As stated above, a process is a system with inputs from vendors and customers, one or more steps to conduct the process, metrics to determine extent of success, outputs moving to a user, and a feedback loop to trigger process alterations. During service design, the steps to the process(es) relating to the proposed service are articulated.

Further, any existing IT service processes on which the proposed will depend are defined in terms of the types of interactions, cost, and service level of support. Services that apply to this definition may include one or more of incident, problem, change, release, configuration, availability, continuity, capacity, and financial management.

Metrics. For both the process(es) and components involved in providing the proposed service, measures and control points are defined. Controls might be human and take the form of checks and balances. For instance, a programmer may make changes to source code but a programming manager must approve the changes in writing. Or, a database administrator may perform data changes but written change approval must be obtained from the client owner of the data. Controls can also be automated, for instance some capability that monitors that all online transactions are actually processed to completion.

The product of service design is a **Service Design Package** which includes sections detailing the following (Rudd & Loyd, 2007):

- Business requirements that are the bases of the design
- Conditions of use – number of users, locations of users, and services used
- Contact information for all stakeholders including business, customer, vendor, IT, etc.

- Functional requirements of the new or changed service including inputs, processes, outputs, quality, security, compliance, privacy, etc. requirements, and deliverables in a signed statement of requirements
- Service level agreement or requirements for revisions to an existing SLA
- Service design and topology/architecture diagrams for IT hardware, network, software, database, etc. resources
- Organizational readiness assessment in terms of user knowledge and skills, financial ability to support the total cost of service ownership, and estimates of development and operational costs (to be refined throughout the build/transition process)
- Service support required including, e.g., service desk
- Transition plan for acquisition of IT resource assets whether guild, buy, outsource, or other
- Plans for quality assuring all service components and acceptance criteria
- Operational acceptance process and criteria for all involved service areas and environments (e.g., cutover plans such as pilot operation and phased transition by function, geography, time, or other).

Tactical Process Design

During the development of the initial service offerings, and as part of continuous improvement, several tactical services are also planned and deployed. This section discusses the definition of five key management processes that underlie IT operations' daily functioning that represent tactical planning in IT: Capacity, availability, continuity, service level, and financial.

Capacity Management Design. Capacity design includes both initial planning for infrastructure and the on-going evaluation of 'sizing' infrastructure to support new functionality. During initial planning the components of capacity, devices, speed, maximum load, etc. are defined and a collective throughout estimate is developed. Usually the manufacturer of the server/computing capacity participates in this activity.

At a minimum, a capacity plan should include the following (Rudd & Loyd, 2007):

- Business scenarios addressed by the plan
- Scope, e.g., a particular change such as move of 1,000 PCs to a new facility
- Methods used – statistical workload and service level forecasts, tools used, etc.
- Assumptions, e.g., capacity drivers
- Service summary
  - Current service levels and resources
  - Forecasts needs
- Resource summary
  - Current utilization
  - Forecast utilization
- Alternatives for resource provision
  - Technology/software alternative
  - Cost and other information on which the decisions are based
- Recommendations

Capacity planning as an on-going activity evaluates the resource requirements of proposed service changes to determine if the change can be absorbed and meet the proposed service level without adversely impacting other service levels. If an impact is expected, testing to determine the extent of impact would be conducted and a request for change and capital expenditure request would be initiated if added computing resources were required.

Periodic review of capacity usage, at a minimum monthly review, is conducted. The review evaluates peaks and trends to predict needed additions to computing resources before any negative impacts to service levels.

Availability Management Design. Availability refers to the cost-effective management of IT operations in full production status that meets business requirements. In the planning phases of the data center and any proposed service, desired availability is defined in terms of number of hours per day and number of days per week computing resource use is anticipated, taking into account the service needs defined in the service catalog. Once in production status, all IT resources are monitored with metrics maintained to prove availability and support capacity planning.

Inputs to availability planning include the service catalog, proposed service requirements, the configuration database, and information from security, vendors, and change processes. The outputs of planning and design are an updated IT resource architecture, updated schematics of physical resource connectivity, etc.

Continuity Management Design. The purpose and goal of **IT service continuity management** is to "support business continuity management process by ensuring that the required IT technical and service facilities (including computer systems, networks,

applications, data repositories, telecommunications, environment, technical support, and service desk) can be resumed within required, and agree, business timescales" (Rudd & Loyd, 2007, p. 262). The first step to continuity design is risk assessment with estimates of timing and severity for different types of computing facility risks.

For instance, the threat of catastrophic outage requiring a move to a backup location might be estimated to occur once every 20 years at a cost of $20 million for three months. A threat of critical device failure might be estimated to occur once per year at a cost of $10,000. These two threats indicate the breadth of consideration for continuity design and provide guidance on how to cost the resources that should be spent in planning, testing, and maintaining a continuity plan. To determine the amount to be spent, a discounted cash flow analysis for the $20 million is performed and averaged, in this example, with the annual $10,000 cost of the device threat. The amount remaining is an estimate of what should be spent on developing and testing continuity plans. Some frequent/low impact risks might be considered 'acceptable' in that they fall under the realm of normal operating procedure for the data center. All other risks undergo risk mitigation to the extent possible and have planning conducted for recovery should they occur.

A continuity plan include items such as:

- How an emergency is decided and announced.
- Succession plan management responsibility list with contact information
- Backup meeting site should the primary site be unavailable
- Plans for obtaining money, plane tickets, living accommodations, backup up tapes (or other storage devices) from off-site storage, and so on.

- Contact information and contract or policy information for all vendors, insurance companies, and contractors.
- The steps to be taken to resume operations at the backup site.
- Addition of services over time should an outage continue.

In planning continuity, risk assessment includes business impact analysis, the organization's business continuity plans, the portfolio of applications and criticality assessment of recovery for each, type of backup/recovery required for each risk, security or compliance issues needing special consideration, service level commitments, vendor contracts (operating level agreements – OLAs). The output is a full IT Service Continuity Plan (ITSCP) that documents all of the mitigation and recovery plans. The goal is that when an emergency or significant outage occurs, few immediate decisions beyond declaring the emergency are needed. The decision making is documented within the plan, which is executed to resume operation as expeditiously as possible.

Service Level Design. Service levels document the agreement between the IT organization and a specific customer about the service to be provided, the level of service guaranteed, remedies for failure to meet guaranteed resource availability. Included in a service level agreement are the following:

- Hours of availability
- Procedures for alternation of available hours including special timing, exceptions or conditions for variations
- Requirements for preventive or other maintenance our 'housekeeping' activities such as backup or report processing

- Target IT resource availability in terms of a percentage of available time within the hours of operation. In contrast to traditional IT availability, this number relates to the 'service' not individual technologies.
- Expected response time for what percent of time (e.g., <2 second response time for 98% of available time). This requires an unambiguous definition of both response and available time.
- Batch turnaround times as appropriate, with details on deliverable inputs and outputs
- Estimated number of outages that do not result in penalties (e.g., two outages per year of more than three hours). This is often expressed as mean time between failures (MTBF).
- Response time to detect and repair outages (e.g., critical applications might have a detection time of five or fewer minutes and an average repair time of 15 minutes). This is often expressed as mean time to repair (MTTR).
- Procedures for emergency or temporary extension to service support
- Contact information for client and IT organization SLA responsibility
- Minimal functionality accepted before the SLA is considered breached
- Requirements for change, continuity, security, privacy, and printing
- Requirements for service reporting and review of the SLA and SLA performance

Financial Management Design. Begun during strategy definition**,** financial estimates for service provision are refined during service design. As design progresses, each component is assigned its actual acquisition and/or disposition cost to provide a more

accurate assessment of the service cost. When design is complete, if there is substantial deviation of actual versus estimated costs beyond some predefined threshold, e.g., ±20%, the CIO may want to report the deviation to the executive committee for further deliberation.

Service Portfolio and Service Catalog Design. The Service Catalog is the official source of all services that are in the operational environment. The catalog may include services being transitioned into production within a short period of time (e.g., a week). For each service, interfaces (triggers, inputs, and outputs), relationships, and supporting services are defined for each service. In addition, any relationship of a given service to the service portfolio is identified.

The Service Portfolio includes the "Service Catalog, a Service Pipeline, and Retired Services" (Nieves & Iqbal, 2007, p. 119). The catalog includes all services that have transitioned into production. The pipeline defines the conduct of continuous improvement for service management, including services that are awaiting development and resource allocation. The Retired Services directory documents all services that have transitioned out of production, have been superceded, or otherwise are no longer performed.

The Service Portfolio is maintained through the strategy processes, but the all three documents are updated throughout design as new services move toward implementation and others move toward retirement. A typical portfolio entry would include information such as a service description, the business case and value proposition, priorities, risks, cost, pricing, and details of the offering (Nieves & Iqbal, 2007). A typical service catalog entry would include for each service, description of the

service including level of service in terms of response time, hours of operation, etc., the software product(s) used, databases involved, acquisition procedures (as needed), support terms, pricing and any chargeback (Nieves & Iqbal, 2007) .

SERVICE TRANSITIONS

Once all design is complete and quality assured, the service begins its transition to operational status. Transitions are initiated either from a change in strategy as discussed above, or from an incident outage, or from problem assessment and change, or from some user or staff request. Transition planning encompasses change management, actual release of the new service into operational status, and updates to the configuration management database (CMDB) from resulting configuration changes.

Change Management.

Changes to IT resources and operating environments account for the majority of incidents and outages. Therefore, managing change is an important activity that must accommodate both normal and emergency situations. A normal, planned change is any change request from users through normal negotiation, service desk requests, or non-outage situation. Emergency changes are changes that result from an outage incident from which recovery requires a change to source code. Planned change should account for 70+% of the changes and should have a target percentage set by the executive board or by the Change Advisory Board (CAB). The CAB is a group of relevant stakeholders from user and IT organizations that approve and oversee change activity.

The Change Manager is an individual who is responsible for planning changes, monitoring the acquisition or development of computing or service assets, working with

the CAB to approve, prioritize, and assess the impact of changes to ensure proper management.

During the design stage for change activities, the process through which changes will be made and the metrics used to report change success are designed. Once designed, change management is one of the principal transition processes to minimize incidents and mitigate the damage of bad code, patches, or resource changes if erroneously moved into production.

The Change Manager is also responsible for communicating the 'forward schedule of change' to all stakeholders for comment and their knowledge, of change status as a change moves through acquisition and/or development, and of the results of change activities, audits, or other reports on change success.

The basic steps to the change process are the following (Conger & Schultze, 2008):

- Receive and log a Request for Change (RFC) from an end user or through incident, problem, configuration, capacity, or other IT management process.
- With the CAB as needed, prioritize the RFC as minor, major, or emergency.
- Assign the work to be accomplished with the appropriate stakeholder/IT resource manager.
- Work with the Release Manager to ensure an acceptable level of quality and that testing of all relevant elements of the change, whether hardware, software, data, or service, are completed successfully.
- Work with the Configuration Manager to define the changes to the CMDB that result from the change.

- Work with the CAB to approve and schedule the change for implementation in the production environment.
- Work with the Release Manager or other technical management to move the change into production and close the RFC.
- Conduct a post-implementation review to ensure that failures in a change result in learning on how to prevent such failures in the future.
- Provide monthly (or other scheduled) reports to all stakeholders, CAB, and change participants to report on the success of change efforts.

Release Management. Like change management, release management is designed during an initial design phase. Elements of design are policy, processes, and communication plans of release management, a 'definitive software library' to contain baseline copies of all production software, and the roles and responsibilities of different constituencies for IT resource planning, testing, and implementation.

Once release management practice begins, the role of the Release Manager is to assure the tested quality of release items and to guarantee that testing of back-out plans provide for a roll-back to the prior version should a release not work in the production environment. The Release Manager works with the Change Manager to develop and maintain the forward schedule of change, and works with the appropriate technical groups to distribute and install the change and any related documentation. Once a change goes live, the CMDB is updated as appropriate for related changes. The known errors database (KEDB) used in incident management (see below) is updated with the problems that are prevented or mitigated as a result of the change. Finally, the Release Manager

participates in the post-implementation review of changes that is conducted by the Change Manager.

Configuration Management. In the mainframe era, maintaining control over assets was relatively simple in that listing a single 'server' and its related components and dumb terminals was straightforward. In the era of 'asset management' information about IT resources maintained included such items as item description, part number, asset id (if one was maintained by the organization), ownership status, contract/lease information, location, cost, depreciation method, and date of installation. Other information that might be maintained on individual items might include preventive maintenance plans and actual performance, downtimes and reasons, and so on. Since all software ran on a single platform (or few platforms), less information was maintained and it might be in the form of 'run books' that were used by operations to initiate, terminate, and prove processing.

In the client-server era, companies moved to having one application per server which simplified monitoring but eventually led to outages of HVAC capabilities and of physical space in data centers.

In the current era of server consolidation, multi-server applications and databases, regional and federated computing, the IT computing environment is significantly more complex with inter-relationships between equipment and network connections that must be known, documented, and monitored to effectively manage on-going operations. In this environment, the purpose of a **configuration management database (CMDB)** is to document and provide management of disparate operating environments. Beginning in a top-down manner, each component of a service is defined as a **configuration item (CI)**. Configuration items are then documented in the CMDB to provide asset identification

information as well as relationship information to all other CIs. Thus, when an outage occurs, a check with the CMDB can immediately provide information on other CIs that will be affected. As appropriate, user stakeholders can be notified of outages before they notice them to minimize transaction processing effects. In the ideal world, the CMDB would include an auto-discovery capability so that any new equipment, software, or changes to existing CIs can be identified and investigated as required. Also ideally, the CMDB would be an integral part of network operating software so that as outages are discovered electronically and reported to an operator and the service desk, all related CIs would also be automatically identified. At present, CMDB management has immature technology that cannot fully realize the goals of configuration management.

Service Knowledge Management. In any complex organization, reasoning, decisions, compliance activities, work flow, and process knowledge are easily lost as people move to new duties in other parts of an organization. The purpose of management service knowledge is to mitigate that loss of knowledge by the creation of a Service Knowledge Management System (SKMS) (MacFarlane, 2007). In defining the SKMS, the cost of capturing and maintaining information about organizational activities must be weighed against the value to be gained from the expenditure. Value to the organization through future use of the information in other parts of the IT organization or by other individuals in the same position that created some information are key to determining which information to collect in the SKMS.

During the design phase, the SKMS data architecture is documented to accommodate designated data and information items, along with the rationale for the inclusion of a information item type. For each item, ownership, access and other

restrictions, privacy and security, and intellectual property information is provided as required (MacFarlane, 2007). Along with the designer of change management, the means for modifying the SKMS through the change process is defined.

Ideal presentation of the SKMS is via a Web portal that includes areas for the following knowledge areas (MacFarlane, 2007, p. 279):

- IT Governance, e.g., reports on governance decisions, continuous improvement projects and status
- Quality management, e.g., policies, processes, checklists
- Training, e.g., courses, trainers, schedules
- Service transparency information, e.g., dashboards of service status, KEDB for lookup of workarounds, solutions to incidents, etc.
- Asset and configuration transparency, e.g., CMDB data
- Service and support, e.g., forward schedule of change
- Self-service capabilities, e.g., to request a change, sign up for a training session, request a report.

Each area of information presented in the SKMS should be considered for "search, browse, store, retrieve, update, publish, subscribe, and collaborate" functionalities (MacFarlane, 2007, p. 279).

SERVICE OPERATIONS

Service Operations encompasses all activities needed to maintain on-going IT Operations on a daily basis. Part of operational management is the monitoring of the computing resources to ensure they are in operational status and that all requests for human intervention are met. Another part of operational management is the recovery from

outages that are part of incident management. The third part of operational management is the monitoring and assessment of metrics for all services to ensure meeting of service level agreements and to provide reporting to all stakeholders. Each of these activities is discussed in this section.

Service Operations Practice. In general, **service operations** is the daily operation of the IT utility such that it satisfies all service level agreements. ITIL v3 explicitly recognizes several organizational functions that were only implied in other versions: IT Operations, Application, and Technology Management. Incidents and problems are managed through the Service Desk while three other new processes are triggered by requests through the Service Desk that are fulfilled by other functions within the IT organization. New to ITIL v3, Request Fulfillment, Event Management, and Access Management acknowledges the diversity of activities performed by Service Desk; that is, the activities are not new, the explicit handling of them by formal processes is. All of these functions and processes are described in this section.

IT Operations Management Function. This function manages and maintains the IT infrastructure on a daily basis. Measures are taken to prove that the IT capabilities delivered have the utility (functioning) and warranty (SLA) agreed upon for each service.

The operations function focus is on cycles to be performed and repeated daily over long periods of time. Successful staffing of operations requires highly skilled technical staff who understand both the service aspects of their positions as well as the technical knowledge to keep the operation running. This external, customer service focus and internal, technology focus are one of the tensions to be balanced in service operations. Another tensions to be balanced include the desire for stability in the

operational environment and the need for fast, flexible response to changing conditions or opportunities. The other tensions to be balanced are quality and cost of services and the need to be proactive with being reactive.

Application Management Function. In pre-service organizations, applications were managed as a separate and distinct area within the IT function. In ITIL v3, application management focuses on the managing of projects to design, development, and deployment of applications, including the management of outsourced development and the acquisition of software as a service, application services, or off the shelf software. Application acquisition specialists coordinate much early acquisition work with specialists performing capacity and availability planning to develop the architectures and infrastructures to accommodate the new functionality. Sub-functions performed in the applications area include the development of regression test packages to be exercised whenever a change to an application is made, testing and quality assurance for application software regardless of its source. While many of the internal processes of applications are separate from most service processes, significant interfaces and overlaps occur in change/release activities and in incident/problem resolution.

Technology Management Function. IT Operations requires considerable technical skills to perform many of the tasks, development of capacity and availability plans, troubleshooting for problems in complex infrastructures, guaranteeing integrity, confidentiality, and availability of data in multi-location, federated infrastructures, and so on. The Technical Management function is responsible for oversight and management of all IT infrastructure including physical plant and equipment as well as the intellectual property relating to planning, troubleshooting, etc. The key sub-processes at the

oversight level relate to selecting and deploying people with the needed skills and expertise in the right positions to leverage their abilities.

Service Desk Function. The Service Desk provides a single point of contact for IT users to handle queries, problems, incidents, requests for changes, requests for fulfillment, and requests for access rights to data and/or applications. Typical input media to provide entry points for service desk requests include telephone, Web interface, automated event reporting, or IT Operations.

The skills required of Service Desk representatives differ from many in an IT Organization in that individuals should be technically savvy but should have excellent 'soft' skills being able to communicate effectively, remain calm in a crisis, and coordinate with many people in often difficult circumstances. Further, the ability to manage a number of outstanding requests successfully through to completion requires the ability to 'time-slice' and multi-task.

The tasks performed in a Service Desk function for a single request include:

- Obtain request and log it as a Service ticket
- Determine priority and impact, escalating as required
- Note the priority and escalation in the Service ticket
- If the request is an incident, coordinate the return to normal operations, accessing a Known Errors Data Base (KEDB), application, IT operations, and other support personnel as required. Monitor return-to-normal time for incidents and escalate as appropriate for the type of incident.

- If the request is for access rights, obtain permissions from the Data or Application Owner and, if given, coordinate through the Access Management process to accomplish the work.
- If the request is for new or changes IT resources, obtain permissions from the individual's manager and, if given, coordinate through the Request Fulfillment process to accomplish the work.
- If the problem is an event, determine with IT Operations if the event has escalated to the status of an incident and, if it has, coordinate through the Event Management process to return to normal operational status.
- For all incidents and problems, update the KEDB with the resolution to speed the recovery should the incident again occur.
- Over time, if trends are noticed in incidents, or if an incident has a significant negative financial impact on the organization, open a problem management request and coordinate with the Problem Manager (which may be the Service Desk representative) to initiate the change process. Coordinate with the Problem, Change, and Release Managers as needed.
- For all requests, communicate status of the open request as appropriate.
- When the incident, request, or problem is complete and the requirements satisfied, log the service request as closed and notify the responsible user and other stakeholders as appropriate.

Incident Management Process. An incident is any "unplanned interruption to an IT Service of a reduction in the quality of an IT service" (OGC, 2007, p. 23) . Incidents are managed by the Service Desk as described above. Each incident moves through a series

of sub-processes to detect, diagnose, repair, recover, and restore normal service, with optional escalation if the recovery exceeds established thresholds of time, number of people affected, etc (Rudd & Loyd, 2007). Upon successful restoration of normal operational status, the work-around or remedy to the incident are documented in the KEDB to reduce recovery time in event of a recurrence.

The goal for incidents is resolution in a single phone call or contact. If the incident is a recurring issue, accurate and timely updating of the KEDB and work-arounds or restoration activities can facilitate this goal. If the incident is novel or not documented, several phone calls, contacts, emails, or other might be required for full resolution.

Problem Management Process. Over time, incidents may evidence trends of some type that are determined to be sufficiently severe or occur frequently that a recommendation for problem diagnosis and remedy is made, usually by a Service Desk representative, an operations staff, an applications staff, or by a specialist who monitors incidents over time.

Once problem remedy is selected, the Service Desk opens a 'problem ticket' and turns the problem over to a Problem Manager. The Problem Manager convenes a team of specialists to perform a root cause analysis to identify all possible and, eventually, the specific trigger of the problem. Upon identifying the problem source, it is fixed and unit tested. The specialist team works with the Release Manager to acceptance test the fix and a back-out solution, thus qualifying the fix for entry in the forward schedule of change. When all approvals and acceptance are obtained, the fix is moved into production on the assigned date and retested to ensure that it will not cause further outages.

When all problem resolution is complete, the Problem Manager coordinates with the Service Desk Representative to close the problem ticket and notify all affected people of the change to the operating environment. The SKMS is updated by the Service Desk to publish problem restoration information.

Request Fulfillment Process. Many requests of the Service Desk are for non-critical, daily activities performed by Service Desk Representatives, for instance "Where is the 'any' key?" Since mundane requests still take time to understand, coordinate, and monitor to completion, they now have a process that recognizes that recurring steps are needed to manage the request. Requests might be for information about a subject, contact information, location of knowledge information, and so on. The SKMS can assist the Service Desk Representative in quickly replying to the request. The goal for most requests is to have them resolved or satisfied in a single phone call.

Event Management Process. An **event** is "a change of state which has significance for the management of a configuration item or IT service" (OGC, 2007, p. 20). State changes most often occur in hardware/software operating environments and may cause an outage or interrupt normal operation, or not. Often events presage an outage and astute management can remedy the problem before an incident occurs.

Events can be monitored by either active or passive tools. Active tools, in essence, ping equipment to determine its status, creating an alert event should status be anything other than normal. Passive tools detect and communicate events, usually reporting through operating system or network monitoring software. Both types of tools can be used in tandem in a single environment. Event response is performed by IT Operations personnel who escalate, as appropriate, to maintain normal operational status.

Monitoring of events is important to identify impending end to a component's operational life, account for time spent troubleshooting and restoring events, and for justifying the levels of expertise needed in IT Operations. Further, event frequency is one measure of overall environmental stability that can be the justification for infrastructure improvements.

Access Management Process. Frequently, requests of the Service Desk are to provide access rights to technology, applications, or data. The purpose of the Access Management Process is to enable access to one or more services or resources while protecting against unauthorized access. Following policies and processes relating to security rights management and accessibility to resources, the Service Desk Representative coordinates the obtaining of needed approvals, documents approvals for compliance as needed, and motivates the provision of access rights.

SUMMARY

This chapter summarizes the IT Infrastructure Library version 3 (ITIL v3) which provides for servitization of the IT operational environment and customer support for its constituents. V3 introduced the service life cycle to better accommodate a philosophy of continuous improvement, and links service initiation to strategic change in the organization. Upon implementation of key services, alignment of IT and its services to the organizational strategy, tactics and specific initiatives is simplified. Overall, ITIL provides for rational, comprehensive management of IT Operations integrating key coordinating activities to other IT organizations. The sum of the services and activities that comprise ITIL ®, can result in management of customer demand from its inception

through its useful life in the organization. This demand management *is* service management.

Reference List

Cartlidge, A., Hanna, A., Rudd, C., Macfarlane, I., Windebank, J., & Rance, S. (2007). *An Introductory Overview of ITIL v3* London, UK: itSMF Ltd.

Conger, S. (2008). *Process Mapping and Management*. NY: Forthcoming.

Conger, S. & Schultze, U. (2008). IT Governance using ITIL, Cobit, and CMMI. Chicago, IL, Society for Information Management, Advanced Practices Council. Ref Type: Unpublished Work

Dorgan, S. and Dowdy, J. (2004) When IT Lifts Productivity, McKinsey Research in Brief.

Farlex (2008). The Free Dictionary. http://www.thefreedictionary.com/service [On-line].

Fitzsimmons, J. A. & Fitzsimmons, M. J. (2005). *Service Management: Operations, Strategy, Information Technology*. (5th ed.) New York, NY: McGraw-Hill/Irwin.

Fleming, W. (2005). *Using Cost of Service to Align IT* Presentation at itSMF-USA, Chicago, IL.

MacFarlane, I. (2007). *ITIL v3 Book 3: Service Transition*. London, UK: The Stationary Office of the Office of General Commerce.

Nieves, M. & Iqbal, M. (2007). *ITIL v3 Book 1: Service Strategy*. (3rd ed.) London, UK: The Stationary Office of the Office of Government Commerce.

OGC (2007). *ITIL V3 Glossary* London, UK: The Stationary Office, The Office of Government Commerce (OGC).

Rudd, C. & Loyd, V. (2007). *ITIL v3 Book 2: Service Design*. London, UK: The Stationary Office of the Office of General Commerce.

TSO (2007). *ITIL® V3 Glossary v01* London: Office of Government Commerce.