ITIL CAPACITY MANAGEMENT DEEP DIVE
CMG 2005 Session 313

Chris Molloy

ITIL™ is continuing to grow in acceptance in IT environments as a model for best practices. This paper provides a low level analysis of the ITIL capacity management discipline. The paper describes the differences between business, service, and resource capacity management, and the need for each to have a proactive capacity management process. The paper will describe the elements needed for an ITIL based capacity plan, a discussion on what several companies have done with ITIL capacity management, and lessons learned from implementing ITIL capacity management discipline in their environment.

1 Background

As the part of the continuous improvement in the computer industry, many organizations have expressed a desire to have a model of information technology. Some organizations have created their own model. Some organizations have outsourced to service providers which have their own proprietary model. An example of this is IBM's™ Information Technology Process Model (ITPM). Still other organizations have looked for more of an open source model.

Information Technology Infrastructure Library (ITIL) is considered such an open source model. While it is maintained by the Office of Governmental Commerce (OGC) of the United Kingdom, the model is available to be used by any organization for modest investment.

The ITIL framework is divided into five principal elements: the business perspective, managing applications, delivery of IT services, support of IT services, and managing the infrastructure. Capacity Management (CM) is considered part of the delivery of IT services, and is described Chapter 6 of the ITIL book, "Best Practice for Service Delivery." [1] For an initial understanding of the ITIL Capacity Management discipline, one should read Chapter 6 which is approximately 40 pages.

The ITIL definition of capacity management is that CM is responsible for ensuring that the capacity of the IT infrastructure matches the evolving demands of the business in the most cost-effective and timely manner. This scope of this definition includes roles and responsibilities previously performed by performance managers and capacity planners. It is both proactive and reactive.

There is continuous work going on to improve ITIL. In April 2005, the OGC published "Results of Public Consultations", which provides recommendations for further ITIL development. The results can be found at http://www.itsmf.com/upload/itil3/consultation_v5.pdf. An overview the current release of the ITIL model entitled "An Introductory Overview of ITIL" is offered in PDF format by the itSMF™ organization. The PDF file is available on the internet. The URL for the file is https://asp5.catalog.com/khamilto2/books/itSMF%20Overview%20Pocket%20Guide.pdf.

This paper is based on the definition and description of capacity management as described in the above mentioned book.

For the last several years, the Computer Measurements Group (CMG) has published papers about ITIL and its role in existing capacity management environments. An example of this was a paper in 2003 entitled "Using ITIL Best Practices to Create a Capacity Management Process" [2]. In 2004, CMG sponsored "The 30th International Conference for the Resource Management and Performance Evaluation of Enterprise Computing Systems." That conference contained a one-day track on ITIL, where the following four papers were presented: "Service Management – The ITIL Framework" [3], "Service Level Management: More Than Just SLAs" [4], "Corporate Performance Management as a Pragmatic Process in an ITIL World" [5], and "ITIL: Make It So" [6].

While there have been several general papers on ITIL, this paper focuses exclusively on the ITIL Capacity Management process. The paper will describe the major sub-processes of ITIL CM: business, service, and resource capacity management. The paper will

then describe the major deliverable from ITIL CM, the capacity plan. To illustrate some of the salient points of the capacity plan, a sample capacity plan is included in Appendix 1. The sample plan is based off of a small office, home office (SOHO) IT environment. While a typical IT environment may be orders of magnitudes larger, the example outlines the necessary elements of the capacity plan that can be scaled to the appropriate level of the IT environment desired. The paper will also discuss what some companies have used ITIL Capacity Management to do to improve their environment.

2 The Sub-Processes of ITIL Capacity Management: Business, Service, and Resource Capacity Management

The purpose of the three sub-processes is to describe what needs to be done within the Capacity Management discipline. It does not say how the processes are to be implemented or by whom. The implementation of a Capacity Management service is an interdependent combination of people, process, and products (tools). ITIL provides the process framework. The CM sub-processes must then be translated into desk procedures which support the sub-processes. The people should be trained to use those procedures, and products should be installed in order to automate some of the labor intensive tasks. There is an inherent tradeoff here between labor and technology. As labor costs continue to increase, it becomes more cost effective to invest in technology to automate many of the tasks needed to support capacity management. An organization must not only decide how much they are willing to invest in each of the capacity management sub-processes, but they must also decide how they are going to split that investment between labor and technology.

Many organizations are looking for products that support ITIL. In order to assist those organizations, a company called Pink Elephant has created an objective set of criteria that meets the minimum functional requirements to support the ITIL framework. They have detailed certification requirements for incident, problem, change, and configuration management ITIL disciplines. Recently, they have added certification requirements for service level management, release management, and availability management. At the time of the writing of this paper, they do not have certification requirements for capacity management toolsets. Additional information can be found on the internet at http://www.pinkelephant.com/consulting/PinkVerify.

The Business Capacity Management (BCM) ITIL CM sub-process is responsible for ensuring that the future business requirements for IT Services are considered, planned, and implemented in a timely fashion. It is the most proactive of the three ITIL CM sub-processes. It is also the most unpredictable as the future is highly suspect to change. It is the least mature of the three sub-processes, and usually has the least amount of investment by IT organizations.

BCM contains the trending, forecasting, modeling, prototyping, sizing, and documenting of future business requirements. These functions require a repository for the numerical information which is called the Capacity Data Base (CDB). Implied in the trending is that future resource requirements have a strong correlation with historical time dependent data. An example of this is email data growth, where growth can be correlated to the growth in number of users (e.g. the trend of an email using 150 Mb of storage on average with the amount of storage needed in the future being computed by forecasting the number of users). In a time of significant merger, acquisition, and divestiture the trending information from both organizations need to be analyzed to determine the new expected trend value for each organization.

The forecasting BCM function includes the translation of new business requirements into changed IT resource requirements. The resource requirements may be increases or decreases. The ITIL CM documentation includes a flowchart of how new requirements for capacity drive the BCM sub-process.

The modeling and prototyping functions of BCM allow the capacity planner to articulate the IT environment in numerical terms, and allow them to perform “what-if” exercises to determine the effects of changes. These models can also be used for sensitivity analysis of potential changes.

The purpose of the Service Capacity Management (SCM) ITIL CM sub-process is to understand service resources, working patterns, peaks and troughs, in order to ensure that the services can and do meet their Service Level Agreement (SLA) targets. Normally these services are mapped to the IT applications (e.g. email, payroll, ordering, etc.) and penalties are assessed for failure to meet the agreement on the quality of service.

SCM contains the monitoring, analyzing, tuning, and reporting of service performance, establishing baselines and profiles of use for services, and managing the demand for services. The SCM monitoring is usually at the application level, and includes application availability, application response, and application usage statistics. An example of this is service monitoring of a web based application for ordering “widgets”, where an SLA could specify that one is capable of ordering widgets 99.9% of the time and that a widget could be ordered in less than two minutes. The SCM analysis function is an application

level view of IT, where one not only understands the application transaction response, but analyzes the critical path of the transaction to determine how to tune the application to improve it. This requires information at both the transaction and sub-transaction level. Reporting of service levels is done on a regular basis to review that service levels are being met, and to calculate any penalties for failures.

The purpose of the Resource Capacity Management (RCM) ITIL CM process is to identify and understand the utilization of each of the components in the IT infrastructure in order to ensure the optimum use of the current hardware and software resources.

RCM contains the monitoring, analyzing, tuning, and reporting of component utilization, establishing baselines and profiles of use of those components. This is the sub-process that is most familiar to people, as people realize that each component has a finite capacity, and that there is the potential for performance problems as resource utilization approaches capacity. Like SCM, it is important in RCM to have real time monitoring, threshold exceptions, and historical reporting, but at a resource level instead of an application level. RCM interacts with ITIL Availability Management discipline, by using load balancing redundant equipment to achieve resiliency. The analysis of this is done as part of a Component Failure Impact Analysis (CFIA).

As the scope of Capacity Management includes hardware, networking equipment, peripherals, and software, RCM should have monitors for each of these components. Many products on the market today monitor both the RCM hardware/software components, and the SCM services (applications).

3 The ITIL Capacity Plan

The ITIL Capacity Plan is listed as the first output of the capacity management process. It is the most important deliverable that goes outside the CM discipline, as it outlines the current environment, and summarizes all the changes that need to be made to respond to the forecasted change in the IT environment. The capacity plan contains all three views of capacity management: the business view, the service view, and the resource view.

It is recommended that the capacity plan be updated at a minimum between each major financial cycle (e.g. yearly) as the recommendations could have significant financial impact.

A sample ITIL capacity plan is included in Appendix 1. This sample was based on a Small Office, Home Office (SOHO) IT environment due to its extremely small size. While the size is small, the concepts apply equally to the larger IT environments. All of the sections of the ITIL Capacity Plan have been included. In order to save space, not all the elements of the plan were completely filled out or filled out to the level of detail that would be expected in a larger multi-user environment. In a large IT environment, some sections will be considerably longer such as the scope section, or the services section. Other sections may be approximately the same length such as the recommendation section in mature IT environment where there are not many recommendations to be made.

The following are comments about the different sections of the ITIL Capacity Plan sample:

1. Introduction (background) – This section allows the capacity planner to state what condition the IT environment is in, and what problems they are facing. It also gives the capacity planner a place to state what recommendations have been implemented from previous plans.

2. Scope of plan – This section lists all the IT resources, and may be better suited as an attachment for organizations with a large number of resources. The scope includes all hardware, networking equipment, peripherals, and software. The collection of this information can be automated with hardware and software inventory tools.

3. Methods used – This section describes where the information came from. The majority of the future information (e.g. new business requirements, forecasts) are expected to come from people. The historical utilization information and modeling information is expected to come from performance tools.

4. Management summary – This may be the first and only section read by many people looking at the capacity plan. It should summarize what needs to be done, when it needs to be done, and how much it should cost.

5. Business scenarios – This section demonstrates the ability of the capacity planner to translate business requirements into IT requirements and then translate those IT requirements to capacity requirements.

6. Service summary – An exhaustive list of all the services (applications) running in an environment can be quite long. The capacity planner should start with the critical applications, either ones that have service levels or those considered for disaster recovery. Additional tools may be needed to determine the resource requirements (profile) at an application level.

7. Resource summary – The gathering of resource

utilization statistics for each device in the environment may be time consuming unless tools are in place to gather this information. Where there is capacity being held in reserve for things like business continuity, workload balancing, or future growth, it should be stated here so that the capacity is reserved and prevented from being reallocated.

8. Options for service improvement – This section should list out all the major options that can be implemented to improve the services.

9. Cost model – This section should outline the cost of implementing each of the options in the previous section. This section is also to include the total cost of running the current IT environment. This could take considerable time to compile, and is constantly changing due to additions and deletions of resources. The capacity planner may choose to compromise the content of this section by including the difference between the costs of the before and after environment instead of computing the total cost of the environment.

10. Recommendations – This section outlines the subset of options that are being recommended for implementation, considering the business environment including cost.

4 Organization ITIL Implementation

This section describes what some IT organizations have done to implement all or parts of the ITIL capacity management discipline.

The first example is an approach to using the ITIL model in a non-ITIL environment. This approach was also capacity management centric, as none of the other ITIL disciplines were involved. As ITIL describes what needs to be done, it can be used as a measuring stick to compare what an existing capacity management implementation contains, and what an enhanced implementation may contain.

An overview of this approach is included in the paper [2] previously referenced. In this case, a spreadsheet was leveraged from the ITIL capacity management discipline, which outlined all the tasks contained in the discipline. The capacity planner performed an evaluation of the current state of capacity management being offered, and a potential desired state. The desired state did not include implementing all the ITIL tasks. Instead, it included a subset of the tasks which represented one possible next state toward a full ITIL CM implementation. This future state was sized, and a proposal made to senior management. In reviewing the situation with the author of the paper, it turns out that this proposal was eventually turned down, in favor of a future state that was in between where the company was, and where they thought they could get to in a reasonable period of time, money, and resources.

This example demonstrates that you can improve an existing capacity management implementation without having to move to ITIL for all your disciplines at once. The items being considered for implementation had little to no interaction with the other ITIL disciplines. This example also demonstrates that this evolutionary process can be repeated several times, allowing a company to progress towards an ITIL implementation at their own pace.

The second example of an approach is more revolutionary in nature. This example is more strategic in nature than the tactical improvements of the first example. The company decided that it wanted to move to an integrated ITIL implementation, of which capacity management was a component. Once that decision was made, a manager was appointed for each of the ITIL disciplines. Each manager appointed a technical representative for their ITIL discipline. All the managers and technical representatives attended formal ITIL education, and became ITIL Foundations certified. The group then met on a regular basis to establish a multi-year roadmap for an integrated ITIL model for their company. The company is in the first year of implementation of this roadmap.

This example demonstrates that a company can also focus on creating a cross discipline integrated ITIL implementation. The company recognized the need for the team to be properly trained in order to be in a better position to create an integrated long term strategy.

The third example of an approach is somewhere in between tactical and strategic. It is an approach taken by a service provider for the accounts that they provide capacity management service for. The provider took the high level ITIL CM process, and took it to a lower level of definition. The provider created a set of roles for personnel that would perform the work, and a set of responsibilities that needed to be performed. The roles were then mapped to the responsibilities using a matrix of which role was to lead the activities and which roles were to participate. There were four roles that were created: the capacity manager, the capacity administrator, the capacity analyst, and the capacity reporter. This service provider had sufficient need for capacity management in order to have multiple people perform each of the multiple roles. A company with fewer requirements may be able to merge some of the roles.

This example demonstrates the need to take the high level ITIL processes, and further define them in a roles and responsibilities matrix. While there may not be as many roles or as many people in a role as in the

service provider case, it is still important that the personnel understand what they are responsible for performing. An example of this is network capacity management. Many organizations have separated the network system support personnel from the other systems programmers. In order to help ensure that all the ITIL tasks are being performed, a matrix between the network and systems areas could be created. The matrix would improve the understanding of who was to perform each of the tasks, including any demarking of responsibilities. This would help ensure that all the tasks were being implemented by one team or the other.

5 Conclusions

The following conclusions can be made from the above discussion on ITIL Capacity Management:

1. Various models exist for defining an Information Technology environment.
2. ITIL is a low cost model to use due to its comparative open source business model.
3. While the nomenclature of capacity management may change between models, the elements needed to provide capacity management service have not significantly changed, and are usually included in each model.
4. The sub-processes of ITIL capacity management provide three views of the same IT environment: the business view, the service view, and the resource view.
5. ITIL provides a framework for what to do, but not how to do it or by whom.
6. ITIL CM service is a interrelationship between people, process and products (tools). Tradeoffs occur between these three items.
7. An organization should decide how much money to invest in capacity management, how much to invest in each of the sub-processes, and how that investment should be balanced between people, process, and products.
8. The ITIL Capacity Plan is a significant output of the ITIL CM process, and should be updated on at least an annual basis.
9. Several organizations have used ITIL CM in different ways.
10. An organization does not have to adopt all of ITIL in order to use the ITIL CM discipline information.

6 References (Acknowledgements)

[1] “Best Practice for Service Delivery”, Office of Government Commerce, ISBN 0-11-330017-4, 2001.

[2] “Using ITIL Best Practices to Create a Capacity Management Process”, C. Molloy, CMG Proceedings, December 2003.

[3] “Service Management – The ITIL Framework”, J Rosenberg, CMG Proceedings, December 2004

[4] “Service Level Management: More Than Just SLAs”, D. Pullorak, CMG Proceedings, December 2004.

[5] “Corporate Performance Management as a Pragmatic Process in an ITIL World”, A. Grummitt, CMG Proceedings, December 2004.

[6] “ITIL: Make It So”, C. Wright, CMG Proceedings, December 2004.

7 Trademarks

Appendix 1 – Sample ITIL Capacity Plan Using Small Office, Home Office (SOHO) IT Infrastructure

1.0 Introduction

This is a capacity plan based on Annex A of the ITIL Capacity Management discipline. A home based IT Infrastructure was used in this example due to small size of the environment. This allows the sample to contain the necessary elements in a few pages. The capacity plan for an actual company would be significantly larger, with each of the sections proportionately longer.

1.1 Background

1.1.1 Organization's current level of capacity

The current organization is thought be in general in an over capacity situation. This is due to the redundancy in the infrastructure to improve resilience, and the over provisioning on the infrastructure to improve performance.

There are two users in the SOHO environment. The portable workstation is the primary workstation for the first user, while the stationary workstation is the primary workstation for the second user. The server is used as the backup for either user in case of a failure of their primary workstation. There is both a wired (100 Mbps) and wireless (54 Mbps) network installed in the environment, with each of the workstations having dual connectivity. This not only provides additional network bandwidth, but allows for workstations to be moved (temporarily or permanently) with minimal impact.

The resources required by the applications in the environment have not significantly increased in proportion to the upgrading of the hardware and software for refresh rates and currency. The refresh rate for the hardware is not on a set schedule, but replaced when there are hardware problems. While this causes an extended unplanned hardware outage, the impact of this outage is minimized by using the server as a backup. The server and stationary workstation are commodity white boxes, built from components. This expedites the repair process by being able to use common parts that are available at local computer stores. The software applications have not changed over the last year. There was a conscious decision to eliminate the need for the swap files on the servers and workstations by installing sufficient memory so as to not overflow into the swap file. This improved performance at a minimal price.

1.1.2 Problems experienced or Envisioned Due to Over or Under Capacity

There are not problems experienced or envisioned due to over capacity. This is occasionally the case where applications grow their requirements over time such as the retention period for email. The first user has a cap on the amount of space to be used for email, while the second user has their email kept on a service provider server (no local copy kept). While there is a large amount of space for file growth, the users are careful enough to only save data that is needed.

There have been a couple of problems experienced this year due to under capacity. While the overall disk space utilization in the environment is low, particular partitions on the server and stationary workstation have come close to capacity. This was caused on the server by installing additional software in the operating system partition, and in the stationary workstation by converting newly bought audio compact disks (CDs) to MP3 format. Removal of unused software was performed on the server to free up space to resolve the first problem, and changing the storage location of new ripped audio files resolved the second problem.

1.1.3 Degree to Which Service Levels Are Being Achieved

While there are no specific service levels for the SOHO environment, both of the users have computer science degrees, and have industry knowledge of what quality of service they feel are appropriate for the work being performed. There were two times this year when the stationary workstation was reported to have "slow" performance. The first problem was due to low signal strength between the workstation on the wireless access point. The wireless access point was moved closer to the workstation, and high gain antennas were installed on both the access point and workstation to resolve the problem. The second problem was due to increased memory requirements for internet browsing. Memory was installed on the workstation in addition to shutting down some unnecessary tasks to decrease memory requirements was implemented to resolve the problem.

1.1.4 What Has Changed Since the Last Issue of the Plan

As this is the initial version of the plan, there are no changes since the last issue of the plan.

1.2 Scope of the Plan

The plan includes the following IT resources:

Name	Purpose	Characteristics
Larry	Server	Hardware - AMD™ Athlon™ XP 1600+ (white box) - 512 Mb memory - 60 GB disk (6 10 Mb Partitions) - CD/RW - DVD/RW - diskette drive (write problems) Software - Windows™ 2000 Server (SP4)
Moe	Portable Workstation	Hardware - IBM ThinkPad T41p - Pentium M 2.00 Ghz - 2 GB RAM - 80 GB disk (All in C drive) Software - Windows XP (SP2) - Rescue and Recovery
Curly	Stationary Workstation	Hardware - AMD Athlon XP 1600+ (white box) - 768 Mb memory - 80 GB (4 20 Mb Partitions) - CD/RW - DV/RW (write problems) - diskette drive - External modem - Video Card Software - Windows 2000 Workstation (SP4) - Legacy DOS Modem Application
PlayStaion/2	Game Console and Media Center	Attached to Surround Sound & TV
Modem 1	Cable Modem	Attached to Time Warner Cable - 10 Mbps
Router 1	Wireless Router	802.11 b/g Wireless Access Point 4 Port 10/100 Router
Switch 1	Network Switch	5 Port 10/100 Switch
KVM 1	Keyboard, Video, and Mouse Switch	4 Port Powered Switch
Printer 1	USB Printer (Color)	Lexmark Z605
Printer 2	USB Printer (Color)	Lexmark Z605
Printer 3	USB Printer (B&W)	Lexmark Z13
Monitor 1	Workstation Monitor	IBM P70 (17 inch)
Monitor 2	Workstation Monitor	IBM P201 (20 inch) – Attached to KVM
UPS 1	Uninterruptible Power Supply	BackUPS 420 Pro
UPS 2	Uninterruptible Power Supply	BackUPS 420

1.3 Methods Used

This section describes the details of how and when information was obtained from the sub-processes. The business forecast was obtained by interviewing the user population in May 2005. The workload forecasts were obtained from the users in June 2005. There are no modeling tools in use to obtain the service level forecasts. As the service level objectives are being met, a decision was made not to install additional tools to collect, store, and report this information beyond what is provided by the particular operating systems themselves (e.g. Windows Task Manager).

2.0 Management Summary

The current IT environment is in excellent health, with minimal to no problems. Proactive investment can be made to prevent problems anticipated by application growth or hardware failure due to lack of refresh. It is recommended that $1220 be invested over the next year for two 160 Mb disk drives (one for the server, one for the stationary workstation), two LCD monitors (one for the server, one for the stationary workstation), and a battery for an existing unused UPS unit for the wireless access point.

3.0 Business Scenarios

There is one expected change to the context of current or envisioned business environment. The users are considering converting their phone service to Voice over Internet Protocol (VoIP). As they will continue to use their current analog phones, the only expected increase in resources will be in the network area. While the local area network should have no problem with the additional bandwidth requirements, the wide are network (cable modem) utilization will be monitored to ensure that this does not become a bottleneck.

There are no moves considered for the SOHO, externally or internally. There are no additional users planned for the environment thru mergers or acquisitions, and there is no decrease in users planned for the environment thru divestitures.

4.0 Service Summary

This section describes a service profile for each service being delivered. In the interest of space for this sample, the majority of the services have not been included.

4.1 Current and Recent Service Provision

Service 1 – email.

Throughput rate – 150 pieces of email a day for first user, 30 pieces of email a day for the second user
Processor usage – Less than 2% CPU utilization when replicating or reading email.
Memory usage – Notes takes 10.820 Mb of memory for first user, browser takes 2.104 Mb for second user
Storage usage – Lotus Notes™ directory 2.51 GB for first user, no additional space for second user
Network usage – 0.01% when replicating for first user on 100 Mbps network, similar for second user

Service 2 – Internet based applications (web browsing using Microsoft Internet Explorer).

Throughput rate – While the amounts of requests are variable, 5 second response or less is expected
Processor usage – Less than 5% CPU utilization for either user
Memory usage – Browser takes 2.104 Mb for either user
Storage usage – Disk storage used considered part of the operating system for either user
Network usage – 2% when active on 100 Mbps network for either user

Service 3 – Music.

Throughput rate – The audio CDs are ripped at 128 kps.
Processor usage – There is less than 5% CPU utilization when recording or playing.
Memory usage – Service takes 4.5 Mb when recording or playing music
Storage usage – Each CD takes about 35 Mb to store music in MP3 format.
Network usage – There is no network traffic when music is played on the stationary workstation, less than 2% when over network.

4.2 Service Forecasts

Service 1 – email. The forecast for email is flat. There is no expected increase in the number of users, and no expected increase in the amount of mail per user.

Service 2 – Internet. Plan for the amount of Internet usage is to double over time as more content is available.

Service 3 – Music. The amount of music acquisition to expected to increase by 4 CDs per month. The forecast for the amount of music playing is flat.

5.0 Resource Summary

This section describes a usage profile for each resource being used.

5.1 Current and Recent Resource Usage

Network resource usage profiles - As the network equipment is redundant, it is significantly oversized, and all the network equipment is running under five percent utilization. The budget for the wired network provides for a 100 Mbps throughput as it comes installed by default on the motherboard with servers or workstations. This bandwidth is designed for a network much larger than the SOHO environment, and is thus overkill. The lowest bandwidth connection is the connection to the Internet thru the cable modem, which has 10 Mbps capacity. Speed tests on this connection indicate a 1.5 Mbps throughput in actual usage.

Computing resource usage profiles

Name	CPU Utilization	Memory Utilization	Disk Utilization
Larry (Server)	Less than 5 percent	86.7%	C Drive – 10.9% D Drive – 0% (spare) E Drive – 0% (spare) F Drive – 82.5% G Drive – 45.5% H Drive – 99.0%
Moe (ThinkPad)	Less than 5 percent	92.2%	C Drive – 55.8 %
Curly (Stationary Workstation)	Less than 5 percent	46.5%	C Drive – 4.4% D Drive – 43.1% E Drive – 89.8% F Drive – 51.4%

KVM resource usage profile – The KVM has only one of the four ports in use, which is for the server. Additional cables are in place to attach the ThinkPad to the KVM in cases where the larger monitor is required.

Printer resource usage profiles – Each printer is only used a couple times a week, so utilization is well within acceptable limits.

Monitors resource usage profiles – As there are only two users, each user can use a different monitor at the same time. While these monitors are essentially connected to a computing resource, remote takeover capability exists on the computing resources should there be a requirement to view the other computing resource while it is in use by another user or in the event of a monitor failure.

UPS resource usage profiles – UPS 1 provides power for the majority of the IT environment (server, cable modem, network switch, KVM, and printer 3). It is currently running at 42% load, with 9 minutes of backup time. The battery was replaced on 9/20/02. UPS 2 provides power for the stationary workstation and its printer, and was replaced in 2004. While the UPS 2 unit does not support usage statistics, it is assumed that the load and minutes of backup time are greater than UPS 1. This assumption is based on the fact that UPS 1 and UPS 2 are the same size, as are the server and stationary workstation (from a power standpoint). As UPS 1 has additional power requirements than UPS 2, UPS 2 is assumed to be running at lower utilization.

5.2 Resource Forecasts

Network forecast – The network forecast for the existing services is expected to be flat. The new VoIP application is

expected to add bandwidth requirements, but it is not expected to be a cause for concern or upgrade.

Computing resource forecast – Growth is expected in the audio service, as it is expected that 4 CDs/month will be purchased, and added to the music collection.

KVM resource forecast – The KVM forecast for the existing services is expected to be flat. No new services are expected to increase KVM resource requirements.

Printer resource forecast – The printer forecast for the existing services is expected to be flat. No new services are expected to increase print resource requirements.

Monitor resource forecast – The monitor forecast for the existing services is expected to be flat. No new services are expected to increase monitor resource requirements.

UPS resource forecast – The UPS forecast for the existing services is expected to be flat. No new services are expected to increase UPS resource requirements.

6.0 Options for Improvement

This section outlines the possible options for improving the effectiveness and efficiency of service delivery.

Network – There is no need to upgrade the network to take advantage of technological advances. The next step would be to go to 1G Ethernet. As the current 100 Mbps network is overkill, there is no need to upgrade.

Computing - As there is a single primary workstation for each of the users with the server as backup, there are no options for merging workloads into fewer computing resources. There are no resource or service tuning plans. There are no plans for rewriting legacy applications. The disk space for the music service (89.8%) is contained at this time by moving new growth to a separate partition (D drive), but the D drive should be monitored to ensure that no additional space is needed on the stationary workstation. The disk space for the primary partition on the server (99.0%) is contained at this time, but should be monitored to ensure that no additional space is needed.

KVM – There are no KVM improvement options being considered at this time.

Printer – There are no printer improvement options being considered at this time.

Monitor – Both monitors are over 8 years old. A monitor refresh should be considered for each monitor. Consideration should also be given to increase the size (19 inch for monitor 1, 21 inch for monitor 2) and improve the technology to LCD thin screen monitors.

UPS – Tests on the UPS 1 battery confirm that it is still effective. It is anticipated that it will need to be replaced in late 2005 or early 2006. The UPS 2 battery should last another several years. The only major piece of IT infrastructure that is not under protection of a UPS is the wireless access point. The wireless access point is not physically located within range of the existing UPS units. An additional UPS should be considered for the wireless access point, as residential power is being used for the SOHO environment, and is not as stable as commercial power.

7.0 Cost Model

The anticipated costs for new equipment are as follows:

1. If the either of the physical disk drives need to be replaced, they should be replaced with 160 GB 7200 RPM drives which cost approximately $86 each.
2. A 19 inch monitor costs approximately $325.
3. A 21 inch monitor costs approximately $700.
4. A battery for a UPS costs $23 (there is an existing UPS available that just needs a battery).

The total equipment cost for all upgrades is $1220. This does not include the labor costs to install the equipment. As both users have computer science degrees, it is expected that they can install the equipment during normal business hours (no additional labor charges). There are no software upgrade charges as a result of this new equipment.

The VoIP converter has already been purchased (it was free with subscription rebates) so it was not included in the cost model.

In a traditional capacity management plan, the current and future cost of providing IT services would be obtained from the IT Financial Plan of the ITIL Financial Management process. It not included in this example. The costs represented above are in addition to the existing cost of service.

8.0 Recommendations

This section contains a summary of the recommendations made in the previous plan, and their status, for example rejected, planned, or implemented. As this is the initial plan, there was no previous plan to comment on. The following information is based on implementing the recommendations indicated by the new cost model.

8.1 Business Benefits to be Expected

The conversion to VoIP is expected to save $25 per month over the existing phone service. In addition, the SOHO will go from 200 minutes of long distance service to unlimited long distance service. Installation of new disk drives will not only provide space for the services needed, but will improve the overall performance of all services that use disk space since the new disk drives have improved performance characteristics over the existing disk drives. The upgrading of the monitors will provide for additional customer satisfaction, as there will be more screen space, and should decrease the amount of scrolling in web based applications. There would also be power savings, space savings, and ergonomic improvements in the newer thin screen monitor technology. The installation of an UPS for the wireless access point will improve network availability during the periods of intermittent residential power variances. Previous variances have caused one of the motherboards of the server to be replaced, so it is worth the investment. Part of this increase in network availability will come from not having to recycle the computing resources in the event of a wireless access point power failure.

8.2 Potential Impact of Carrying Out the Recommendations

Except for the investment costs of carrying out the recommendations, all of the potential impact is considered positive to the users. The additional disk space will also provide for the potential for unanticipated growth to existing or other new services. The new monitors have a higher resolution than the current monitors. This provides for sharper graphics such as pictures or web graphics, and the ability to provide more information on the monitor by increasing the screen area resolution.

8.3 Risks Involved

The risk involved in implementing the recommendations is the potential for hardware failure as a result of installation of the new components. This disk installation risk has been evaluated and considered minimal, as disk upgrades have been installed on these systems before with no impact. Replacing a monitor is considered low risk.

8.4 Resources Required

The cost model indicates the hardware resources required. There is no additional software resource required. The labor resources required will be minimal to install the new components.

8.5 Setup and On-going Costs

The setup costs are expected to be minimal since these are performed in house. The on-going costs of the recommendations is the additional space needed to perform backups on these systems (since there will be more data to back up) and the cost of replacing the UPS battery every five years. As backups are currently stored on DVD, the additional DVD cost is expected to be $10/year. The additional battery cost is expected to be $23 every five years.