


# Cloud Computing Specialist Certification Kit

## Virtualization



The Art of Service

# Cloud Computing Virtualization Specialist Complete Certification Kit:

## Study Guide Book and Online Course

**Notice of Rights:** Copyright © The Art Of Service. All rights reserved. No part of this book may be reproduced or transmitted in any form by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher.

**Notice of Liability:** The information in this book is distributed on an “As Is” basis without warranty. While every precaution has been taken in the preparation of the book, neither the author nor the publisher shall have any liability to any person or entity with respect to any loss or damage caused or alleged to be caused directly or indirectly by the instructions contained in this book or by the products described in it.

**Trademarks:** Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and the publisher was aware of a trademark claim, the designations appear as requested by the owner of the trademark. All other product names and services identified throughout this book are used in editorial fashion only and for the benefit of such companies with no intention of infringement of the trademark. No such use, or the use of any trade name, is intended to convey endorsement or other affiliation with this book. ITIL® is a Registered Community Trade Mark of OGC (Office of Government Commerce, London, UK), and is Registered in the U.S. Patent and Trademark Office.

## **Write a Review and Receive a Bonus Emereo eBook of Your Choice**

### ***Up to \$99 RRP – Absolutely Free***

*If you recently bought this book we would love to hear from you – submit a review of this title and you'll receive an additional free ebook of your choice from our catalog at <http://www.emereo.org>.*

### ***How Does it Work?***

*Submit your review of this title via the online store where you purchased it. For example, to post a review on Amazon, just log in to your account and click on the 'Create Your Own Review' button (under 'Customer Reviews') on the relevant product page (you'll find plenty of example product reviews on Amazon). If you purchased from a different online store, simply follow their procedures.*

### ***What Happens When I Submit my Review?***

*Once you have submitted your review, send us an email via [review@emereo.org](mailto:review@emereo.org), and include a link to your review and a link to the free eBook you'd like as our thank-you (from <http://www.emereo.org> – choose any book you like from the catalog, up to \$99 RRP). You will then receive a reply email back from us, complete with your bonus ebook download link. It's that simple!*



## Foreword

*As an education and training organization within the IT Service Management (ITSM) industry, we have watched with enthusiasm as cloud computing and Virtualization have evolved over the past years. The opportunities provided through virtualization have allowed for significant growth within an industry that continues to mature and develop at a rapid pace.*

*Our primary goal is to provide the quality education and support materials needed to enable the understanding and application of virtualization within a wide range of contexts.*

*This comprehensive book is designed to complement the in-depth eLearn Virtualization Specialist program provided by The Art of Service. The interactive eLearn course uses a combination of narrated PowerPoint presentations with flat text supplements and multiple-choice assessments which will ultimately prepare you for the Virtualization Specialist Level certification exam.*

*We hope you find this book to be a useful tool in your educational library and wish you well in your IT Service Management career!*

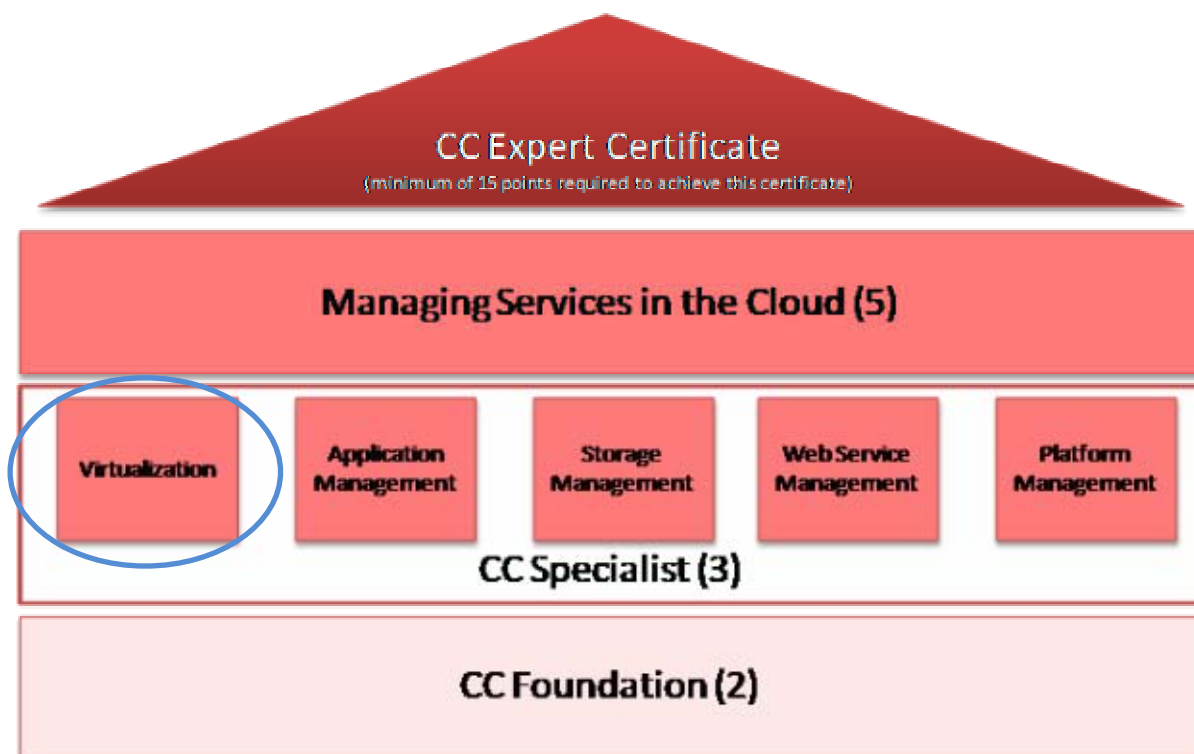
*The Art of Service*

*<http://www.theartofservice.com/>*

## How does the Certification Kit work?

Welcome to the Virtualization Specialist Level Complete Certification Kit. This book is part of the series of books around Cloud Computing and managing the services that are involved in, or utilize, Cloud Computing.

The certification kits are in line with the Cloud Computing Certification Scheme.



After you've read this book, studied the eLearning materials and successfully passed your exam, you can continue with your qualifications through the specialist programs and work toward your Cloud Computing Expert Certification.

In addition to the certification kits, The Art of Service has also developed a series of books on the subject of Cloud Computing. All books are available as e-book, PDF download, audio-book and paperback.

## eLearn Component

This certification kit comes with FREE access to the eLearning program. The following page explains how to access the program materials online.

## The Virtualization Specialist Level Exam

Chapter 8 of this book provides more detail about the certification exam and what the requirements are to pass.

## **How to access the associated Virtualization Special Level eLearning Program:**

1. Direct your browser to: [www.theartofservice.org](http://www.theartofservice.org)
2. Click 'login' (found at the top right of the page)
3. Click 'Create New Account'. If you already have an existing account, please move on to step 5.
4. Follow the instructions to create a new account. You will need a valid email address to confirm your account creation. If you do not receive the confirmation email check that it has not been automatically moved to a Junk Mail or Spam folder.
5. Once your account has been confirmed, email your User-ID for your new account to [ccvirt@theartofservice.com](mailto:ccvirt@theartofservice.com).
6. We will add your account to the Virtualization eLearning Program and let you know how to access the program from now on.

## **Minimum system requirements for accessing the eLearning Program:**

<b>Processor</b>	Pentium III (600 MHz) or higher
<b>RAM</b>	128MB (256MB recommended)
<b>OS</b>	Windows 98, NT, 2000, ME, XP, 2003, Mac OSX
<b>Browser</b>	Internet Explorer 5.x or higher (Cookies and JavaScript Enabled), Safari
<b>Plug-Ins</b>	Macromedia Flash Player 8
<b>Other Hardware</b>	16-bit Sound Card, Mouse, Speakers or headphones
<b>Display Settings</b>	1024x768 pixels
<b>Internet Connection</b>	Due to multimedia content of the site, a minimum connection speed of 256kbs is recommended. If you are behind a firewall and face problems accessing the course or the learning portal, please contact your network administrator for help

## **If you are experiencing difficulties with the Flash Presentations within the eLearning Programs please make sure that:**

- 1) You have the latest version of Flash Player installed, by visiting [www.adobe.com](http://www.adobe.com) and following the links.
- 2) You check that your security settings in your web browser don't prevent these flash modules playing. There is support for these issues on the Adobe webpage.

## Contents

<b><i>FOREWORD</i></b>	<b><i>1</i></b>
<b><i>1 INTRODUCTION</i></b>	<b><i>8</i></b>
1.1 WHAT IS VIRTUALIZATION?	8
1.2 OBJECTIVES OF VIRTUALIZATION	9
1.3 HISTORY OF VIRTUALIZATION	11
1.4 BENEFITS OF VIRTUALIZED TECHNOLOGY	12
1.5 THE VIRTUAL SERVICE DESK	15
1.6 WHAT CAN BE VIRTUALIZED?	15
1.7 RELATED FORMS OF COMPUTING	17
1.7.1 CLOUD COMPUTING	17
1.7.2 SOFTWARE AS A SERVICE - SAAS	18
1.7.3 GRID COMPUTING	18
1.7.4 UTILITY COMPUTING	19
1.8 VIRTUALIZATION PROCESSES	19
1.9 INTRODUCTION REVIEW QUESTIONS	20
<b><i>2 COMMON TERMINOLOGY</i></b>	<b><i>21</i></b>
<b><i>3 VIRTUALIZATION TECHNOLOGIES</i></b>	<b><i>26</i></b>
3.1 UBUNTU (SERVER EDITION)	26
3.2 ALTIRIS	27
3.3 WINDOWS SERVER	28
3.4 SOFTWARE VIRTUALIZATION	29
3.5 VMWARE	30
3.6 INTEL VIRTUALIZATION	31
3.7 RED HAT VIRTUALIZATION	32
3.8 SOFTGRID APPLICATION	33
3.9 LINUX VIRTUALIZATION	35
3.10 DESKTOP VIRTUALIZATION	37
3.11 HARDWARE VIRTUALIZATION	38
3.12 RESOURCE VIRTUALIZATION	39
3.13 PROCESSOR VIRTUALIZATION	40
3.14 APPLICATION VIRTUALIZATION	41
3.15 STORAGE VIRTUALIZATION	42
3.16 VIRTUALIZATION DENSITY	43
3.17 PARA-VIRTUALIZATION	44
3.18 OS VIRTUALIZATION	44
3.19 VIRTUALIZATION SOFTWARE	46
3.20 DATA STORAGE VIRTUALIZATION	58

3.21	INTEL VIRTUALIZATION TECHNOLOGY	60
3.22	THINSTALL VIRTUALIZATION SUITE	62
3.23	NET FRAMEWORK VIRTUALIZATION	64
3.24	WINDOWS VIRTUALIZATION ON FEDORA	66
3.25	STORAGE VIRTUALIZATION TECHNOLOGIES	67
3.26	VIRTUALIZATION LEVEL	69
3.27	SECURITY MONITORING AND VIRTUALIZATION	70
3.28	ORACLE VIRTUALIZATION	74
3.29	VIRTUALIZATION TECHNOLOGIES REVIEW QUESTIONS	75
<b>4</b>	<b><i>ACCOMPLISHING VIRTUALIZATION</i></b>	<b>76</b>
4.1	MIGRATING TO A VIRTUALIZED ENVIRONMENT	76
4.1.1	THINGS TO DO AND CONSIDER BEFORE MIGRATION	76
4.2	THINGS TO DO AFTER MIGRATION	79
4.2.1	A SIMPLE HOW-TO GUIDE	80
4.2.2	FURTHER MIGRATION CONSIDERATIONS	83
4.2.3	RISKS ASSOCIATED WITH VIRTUALIZATION	90
4.2.4	PROBLEMS ASSOCIATED WITH VIRTUALIZATION	95
<b>5</b>	<b><i>MANAGING A VIRTUALIZED ENVIRONMENT</i></b>	<b>97</b>
5.1	SUPPORT ISSUES	97
5.2	MEASURING CAPACITY AND PERFORMANCE	99
5.3	CONTRACTS AND AGREEMENTS SUCH AS LICENSING	101
5.4	ORGANIZATIONAL CONSIDERATIONS	103
	SERVICE MANAGEMENT PROCESSES	105
	IT FINANCIAL MANAGEMENT	105
5.4.1	IT FINANCIAL MANAGEMENT AND VIRTUALIZATION	106
	INFORMATION SECURITY MANAGEMENT	107
5.4.2	INFORMATION SECURITY MANAGEMENT AND VIRTUALIZATION	107
5.5	RELEASE & DEPLOYMENT MANAGEMENT	107
5.5.1	GOALS AND OBJECTIVES	108
5.5.2	SCOPE	108
5.5.3	BENEFITS	108
5.5.4	TERMINOLOGY	109
5.5.5	TRIGGERS AND INTERFACES	112
5.5.6	RELEASE DESIGN OPTIONS AND CONSIDERATIONS	113
5.5.7	RELEASE POLICY	115
5.5.8	RELEASE AND DEPLOYMENT ACTIVITIES	115
	KNOWLEDGE MANAGEMENT	125
5.5.9	POLICIES AND PRINCIPLES OF KNOWLEDGE MANAGEMENT	126
5.5.10	THE SERVICE KNOWLEDGE MANAGEMENT SYSTEM (SKMS)	127
	INCIDENT MANAGEMENT	128
5.5.11	INCIDENT MANAGEMENT AND VIRTUALIZATION	128
5.6	CHANGE MANAGEMENT	129
5.6.1	GOALS AND OBJECTIVES	130
5.6.2	SCOPE	130



5.6.3	DESIGNING AND PLANNING	131
5.6.4	CHANGE MANAGEMENT POLICIES	132
5.6.5	CHANGE MODELS	132
5.6.6	TRIGGERS AND INTERFACES	134
5.6.7	CHANGE MANAGEMENT ACTIVITIES	136
5.6.8	ROLES AND RESPONSIBILITIES WITHIN CHANGE MANAGEMENT	141
5.6.9	KEY PERFORMANCE INDICATORS (KPIs) OF CHANGE MANAGEMENT	142
5.6.10	CHALLENGES AFFECTING CHANGE MANAGEMENT	143
5.6.11	RELATIONSHIP WITH PROJECT MANAGEMENT	144
5.6.12	TYPICAL CONTENTS OF CHANGE DOCUMENTATION	145
5.7	SERVICE DESK	149
5.7.1	GOAL AND OBJECTIVES	149
5.7.2	BENEFITS	150
5.7.3	SERVICE DESK ORGANIZATIONAL STRUCTURES	150
5.7.4	SERVICE DESK TYPES (SKILL LEVELS)	155
5.7.5	SERVICE DESK STAFFING	155
5.7.6	KEY PERFORMANCE INDICATORS (KPIs) FOR THE SERVICE DESK	157
5.7.7	OUTSOURCING THE SERVICE DESK	157

## **6      VIRTUALIZATION AND STORAGE MANAGEMENT      158**

6.1	THE HEART OF CLOUD COMPUTING: VIRTUALIZATION	158
6.2	DEFINING VIRTUALIZATION	159
6.3	WHY VIRTUALIZE?	159
6.4	WHAT CAN BE VIRTUALIZED?	160
6.5	WHERE DOES VIRTUALIZATION HAPPEN?	161
6.6	HOW DOES VIRTUALIZATION HAPPEN?	162
6.7	ON THE ROAD TO STORAGE VIRTUALIZATION	163
6.8	IMPROVING AVAILABILITY USING VIRTUALIZATION	164
6.9	IMPROVING PERFORMANCE THROUGH VIRTUALIZATION	164
6.10	IMPROVING CAPACITY THROUGH VIRTUALIZATION	165
6.11	BUSINESS VALUE FOR VIRTUALIZATION	166

## **7      APPLYING CLOUD COMPUTING TO PROJECT MANAGEMENT      167**

7.1	BENEFITING FROM CLOUD COMPUTING	167
7.2	THE EASE OF LINKING (HYPERLINKS)	168
7.3	SUBSCRIBING TO SUCCESS (BLOGGING)	169
7.3.1	EVERYONE IS A PROJECT MANAGER (OPEN SOURCE)	171
7.4	TREATING THE PROJECT AS PARTS, NOT THE WHOLE (REUSE)	173
7.5	TESTING THE LIMITS (PORTABILITY)	174
7.6	PROCESS REVIEW QUESTIONS	175

## **ANSWERS      177**

7.6	PROCESS REVIEW	177
-----	----------------	-----

## **CERTIFICATION      177**

7.7	CLOUD COMPUTING CERTIFICATION PATHWAYS	177
7.8	ITIL® CERTIFICATION PATHWAYS	180
7.9	ISO/IEC 20000 CERTIFICATION PATHWAYS	181
<b><u>8</u></b>	<b><u>VIRTUALIZATION SPECIALIST EXAM TIPS</u></b>	<b><u>182</u></b>
<b><u>9</u></b>	<b><u>REFERENCES</u></b>	<b><u>183</u></b>

# 1 Introduction

---

## 1.1 What is Virtualization?

One of the most important ideas behind cloud computing is scalability, and the key technology that makes that possible is virtualization.

Virtualization, in its broadest sense, is the emulation of one or more workstations/servers within a single physical computer. Put simply, virtualization is the emulation of hardware within a software platform. This allows a single computer to take on the role of multiple computers. This type of virtualization is often referred to as full virtualization, allowing one physical computer to share its resources across a multitude of environments. This means that a single computer can essentially take the role of multiple computers.


However, virtualization is not limited to the simulation of entire machines. There are many different types of virtualization, each for varying purposes. One of these is in use by almost all modern machines today and is referred to as virtual memory. Although the physical locations of data may be scattered across a computer's RAM and Hard Drive, the process of virtual memory makes it appear that the data is stored contiguously and in order. RAID (Redundant Array of Independent Disks) is also a form of virtualization along with disk partitioning, processor virtualization and many other virtualization techniques.

Virtualization allows the simulation of hardware via software. For this to occur, some type of virtualization software is required on a physical machine. The most well known virtualization software in use today is VMware. VMware will simulate the hardware resources of an x86 based computer, to create a fully functional virtual machine. An operating system and associated applications can then be installed on this virtual machine, just as would be done on a physical machine. Multiple virtual machines can be installed on a single physical machine, as separate entities. This eliminates any interference between the machines, each operating separately.

Although virtualization technology has been around for many years, it is only now beginning to be fully deployed. One of the reasons for this is the increase in processing power and advances in hardware technology. As the benefits of virtualization are realised, we can observe the benefits to a wide range of users, from IT professionals, to large businesses and government organizations.

## 1.2 Objectives of Virtualization

*Before virtualization:*



2x purchase  
2x maintenance  
2x depreciation  
2x floor space  
2x energy usage

There are four main objectives to virtualization, demonstrating the value offered to organizations:

- Increased use of hardware resources;
- Reduced management and resource costs;
- Improved business flexibility; and
- Improved security and reduced downtime.

### Increased use of Hardware Resources

With improvements in technology, typical server hardware resources are not being used to their full capacity. On average, only 5-15% of hardware resources are being utilized. One of the goals of virtualization is to resolve this problem. By allowing a physical server to run virtualization software, a server's resources are used much more efficiently. This can greatly reduce both management and operating costs. For example, if an organization used 5 different servers for 5 different services, instead of having 5 physical servers, these servers could be run on a single physical server operating as virtual servers.

## **Reduced Management and Resource Costs**

Due to the sheer number of physical servers/workstations in use today, most organizations have to deal with issues such as space, power and cooling. Not only is this bad for the environment but, due to the increase in power demands, the construction of more buildings etc is also very costly for businesses. Using a virtualized infrastructure, businesses can save large amounts of money because they require far fewer physical machines.

## **Improved Business Flexibility**

Whenever a business needs to expand its number of workstations or servers, it is often a lengthy and costly process. An organisation first has to make room for the physical location of the machines. The new machines then have to be ordered in, setup, etc. This is a time consuming process and wastes a business's resources both directly and indirectly.


Virtual machines can be easily setup. There are no additional hardware costs, no need for extra physical space and no need to wait around. Virtual machine management software also makes it easier for administrators to setup virtual machines and control access to particular resources, etc.

## **Improved Security and Reduced Downtime**

When a physical machine fails, usually all of its software content becomes inaccessible. All the content of that machine becomes unavailable and there is often some downtime to go along with this, until the problem is fixed. Virtual machines are separate entities from one another. Therefore if one of them fails or has a virus, they are completely isolated from all the other software on that physical machine, including other virtual machines. This greatly increases security, because problems can be contained.

Another great advantage of virtual machines is that they are not hardware dependent. What this means is that if a server fails due to a hardware fault, the virtual machines stored on that particular server can be migrated to another server. Functionality can then resume as though nothing has happened, even though the original server may no longer be working.

after virtualization:



- 1 x purchase
  - 1 x maintenance
  - 1 x depreciation
  - 1 x floor space
  - 1 x energy usage
- => better utilization

✓

### 1.3 History of Virtualization

Virtualization technology has been around for longer than most people realize. It is only now starting to be widely used because of the massive increase in hardware resources.

The concept of virtualization was first devised in the 1960s. It was then implemented by IBM to help split large mainframe machines into separate 'virtual machines'. The reason why this was done was to maximise their available mainframe computers efficiency. Before virtualization was introduced, a mainframe could only work on one process at a time, becoming a waste of resources. Virtualization was introduced to solve this problem. It worked by splitting up a mainframe machine's hardware resources into separate entities. Due to this fact, a single physical mainframe machine could now run multiple applications and processes at the same time.

As x86 became the dominant instruction set architecture in computer technology during the 1980s, and the client-server model was established to allow for distributed computing, the need for virtualization was no longer really required. This is because the client-server model allowed administrators to connect together many low cost workstations. Resources could then be distributed among these workstations using a few powerful servers. The massive use of Windows and Linux based operating systems during the 1990s further solidified the x86 architecture and client-server model, as the dominant model in computer technology.





However, no one could have imagined the massive growth in the use of computer technology, and this created new IT infrastructure demands as well as problems. Some of these problems included:

- Low hardware infrastructure utilization;
- Rising physical infrastructure costs;
- Rising IT management costs;
- Insufficient disaster protection; and
- High maintenance and costly end-user desktops.

The most viable solution to resolve the above-mentioned issues was hardware virtualization and thus, in 1999, VMware introduced their first virtualization application for x86 based systems. Modern day machines can now have their powerful hardware resources split up, just like mainframe machines did during the 1960s, to allow for more efficient usage of processing power and hardware resources.

## 1.4 Benefits of Virtualized Technology

Virtualized technology offers many benefits to organizations looking to migrate from a physical environment to a virtual setting. Some of the specific benefits to businesses are described below:

**Easier Manageability** – Through virtualization, administrators can monitor and manage entire groups of servers/workstations from a single physical machine.

**Elimination of Compatibility Issues** – In the past, running MAC OS, Linux or Windows on the same machine created many compatibility issues. These days, using virtual machines, many different operating systems and applications can run on a single physical machine, without affecting one another.

**Fault Isolation** – Any kind of error within a virtual machine will not affect any other virtual machine. Problems are automatically isolated, which can then be fixed or looked into by an administrator, while all other systems and services continue normal operation.

**Increased Security** – Administrators can separate information and applications on a single physical machine into different virtual machines. This prevents users from being able to access or view what they should not be able to. Furthermore, if a virus gets on to a physical machine, it will usually affect the entire machine. Virtual machines are separate entities, therefore any viruses or problems will be isolated within that virtual machine. The physical machine and all other virtual machines will not be affected. Virtualization also makes it easier to revert back to a previous state. For example an entire virtual machine could be backed up manually at regular intervals or using the virtual machines built in checkpoint feature, it could be reverted to a previous fully working state.

**Efficient use of Resources** – Many virtual machines can run on a single physical machine, utilizing that physical machine's resources much more efficiently than if it was just running a single service or application. If a business previously used five physical servers, to provide five different services, they would now only require a single physical server. This one server would be much more efficiently used and its resources would not be wasted.

**Portability** – Virtual machine data is stored in files on a physical machine. This means that virtual machines can be transferred effortlessly from one physical machine to another, without any changes to functionality.

**Problem-Free Testing** – One or more virtual machines can be set up as test machines. These can then be used to test the stability of certain applications or programs, without affecting the functionality of day to day business operations.

**Rapid Deployment** – The Hard Drive of a virtual machine is often represented as a single file on a physical machine. This means that this Hard Drive can easily be duplicated or transferred to other physical machines. By using one virtual machine as a 'template' machine, its virtual Hard Drive file can be used to rapidly create new virtual machine clones. The advantage of this is that an administrator would only have to carry out an OS installation once.

**Reduced Costs** – Costs are reduced in the form of less physical hardware, less power and cooling requirements, less physical space and less staffing requirements. Once a business has moved to virtualized environment, they begin to reap the many rewards that this kind of solution provides. Less physical machines, means less physical space is required, therefore a business office and building costs are greatly reduced. Both heat and power outputs are also significantly reduced and again this helps to reduce the running costs for a business. Networking costs are also reduced because fewer switches, hubs and wiring closets are required. As you can see, one of the greatest reasons why a business would want to adopt virtualization is simply because of the large amounts of money that they could saved, both in the long- and short-term.

**The Ability to Separate Applications** – Services and applications that may conflict with one another can be installed and run on separate virtual machines. Because these services and applications are still running on the same physical machine, resources and processing power are not wasted.


**Easier Manageability and Improved Uptime** – Virtualization makes administering and managing an IT environment much easier. Less physical machines mean that administrators have less hardware faults or issues to deal with. This means that they can use their time on much more productive and beneficial tasks. Multiple virtual machines can also be configured and monitored from a single location. Again this makes things much easier for administrators, as they do not have to go from physical machine to machine when configuring updates or making changes.

Backing up and restoring a virtualized environment is also much easier than a physical one. If there are problems on one physical machine, an administrator can easily copy over entire virtual machines to another physical machine. A business can then continue operating with a minimum of downtime.

Planned downtime can also be implemented with a minimum of disruption. Again, all an administrator has to do is copy over their virtual machines to another physical machine and then these virtual machines can continue operating as though nothing has happened.

Maximum uptime for modern day businesses is very important, even the smallest amount of downtime can cause a business to lose profits as well as their credibility with clients. Virtualization is one of the ways that a business can ensure they will be running smoothly with an absolute minimum of downtime.

## 1.5 The Virtual Service Desk



Many large organisations operate a virtual service desk whereby the service desk analysts are connected to the user groups and support groups via an internet connection or other network capability. The user groups or end users who call the service desk will not know where the analyst is, however as long as the service desk analyst has access to the knowledge base and the service desk tool to log and record all actions, it shouldn't matter.

## 1.6 What can be Virtualized?

The SNIA taxonomy lists five different types of storage virtualization: disk, tape (media and drive), file system, file, and block.

Disk virtualization is one of the oldest forms of storage virtualization. The physical form of a magnetic disk is a compilation of cylinders, heads, and sectors. Each disk is different based on the numbers of cylinders, heads, and sectors; which changes the capacity of the disk. In order to read or write using the disk, some form of addressing is required. To complete this addressing the exact physical property of every magnetic disk would have to be known – an impossibility. Virtualization is completed by the disks' firmware to transform the addresses into logical blocks for use by applications and operating systems, called logical block addressing (LBA). As magnetic disks are used, some blocks may not be able to store and retrieve data reliably. Disk virtualization allows the magnetic disk to appear defect-free, releasing the operating system to focus on other processes.

When data is read or written to physical tape media, the transmission can fluctuate due to the network or busy clients on the network. The ideal situation is to have a single stream of data. When fluctuations occur, the tape drive reverts to a start/stop mode where a small amount of data is written, the tape is stopped and rewound before another small amount of data is written. This mode increases recording time and wear on the tape media.

Variances like this result in only 15-30 percent of tape media capacity to be utilized. Tape media virtualization uses online disk storage as a cache to emulate reading and writing. This disk acts as a buffer to eliminate the fluctuations and variances present when writing directly to the tape media. Tape media virtualization also emulates large numbers of small volumes while keeping the data on disk. The data is then written to tape in a streaming pattern.

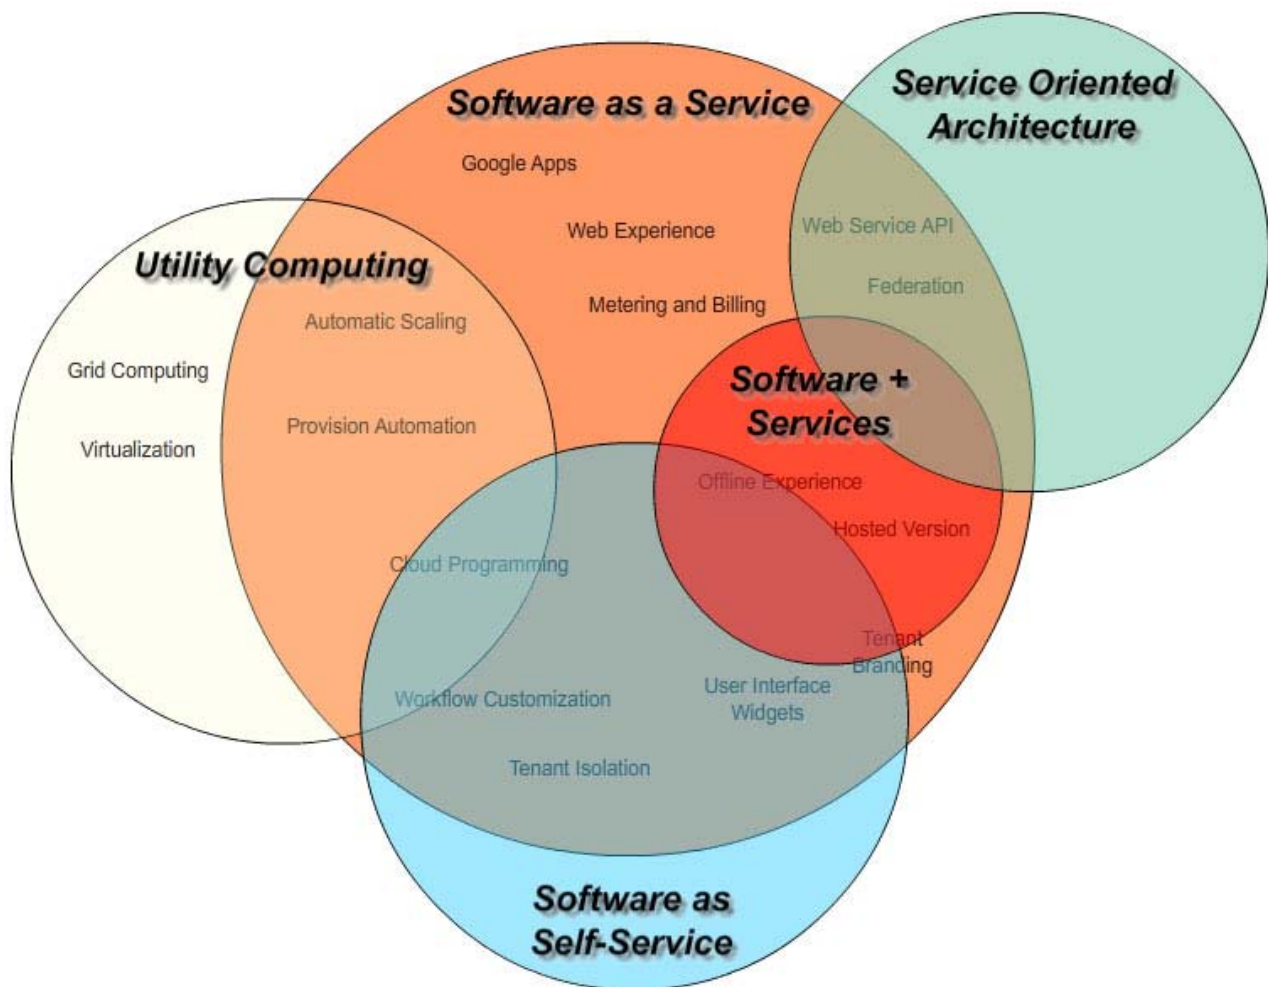
Tape drives used to be dedicated to single servers and then connected to multiple servers. Now tape drives are connected to a server and the activities involving these devices are performed across the network. Unfortunately, a potential risk is two servers writing to the same tape drive at the same time. Virtualized tape drives can make a single drive appear as several virtual drives which are assigned to individual servers as dedicated resources. When a request is made, the virtualization software (tape broker) reserves and maps a physical tape drive to the virtual drive and completes the operations. The physical drive is then placed back into the environment to be used by another virtual drive. Another virtualization form is the use of RAIL (Redundant Array of Independent Libraries). This technique allows several tape libraries to be emulated and distributed as a single virtual library for use by applications.

File systems virtualization provides file systems to multiple clients regardless of the operating systems on those clients. Using a networked remote file system such as NFS or CIFS, this abstraction demonstrates the most important feature of storage virtualization: location transparency. Another form of file system virtualization assists in database management. Most data for databases are located on raw disk drives to maximize performance but are cumbersome to manage. Putting this data into a file system makes the data easier to manage but causes problems in performance. Database virtualization combines the best of both worlds.

Hierarchical Storage Management (HSM) identifies rarely used data and migrates it to inexpensive secondary storage. From the client perspective, the file still exists and there is no need to know the physical location of the file. This is a form of file virtualization.

Block virtualization, or block aggregation, creates a single logical device out of several physical devices. The intent is to have applications see a new virtual disk with a larger range of block addresses; essentially tricking the application into thinking it has more memory than it truly has. The storage customer is concerned with three things: capacity, performance, and availability. When the customer demands more, functions of block virtualization can fulfill these requirements easily and cost-effectively.

## 1.7 Related Forms of Computing



To view this picture as a moving GIF, go to

[http://theartofservice.net/index.php?option=com\\_content&task=view&id=142](http://theartofservice.net/index.php?option=com_content&task=view&id=142)

### 1.7.1 Cloud Computing

Cloud computing refers to the processing and storage of data through the Internet. Computing and storage become 'services' rather than physical resources. Files and other



data can be stored in the cloud and be accessed from any Internet connection. It is a style of computing where IT-related capabilities are provided 'as a service', allowing users to access technology-enabled services from the Internet, or cloud, without knowledge of, expertise with, or control over the technology infrastructure that supports them.

### **1.7.2 Software as a Service - SaaS**

Software as a Service (or SaaS) is a closely related form of computing to the cloud. As with cloud computing, SaaS customers tap into computing resources off-site that are hosted by another company. However, the difference is scale. Cloud computing platforms may combine thousands of computers and storage networks for backup, while some SaaS applications may operate via a connection between only two or three computers.

This means that most SaaS solutions fall under the larger cloud computing definition. However, there are a few distinguishing features of SaaS that separate these solutions from cloud computing. SaaS is software that is owned, delivered and managed remotely by one or more providers. It also allows sharing of application processing and storage resources in a one-to-many environment on a pay-for-use basis, or as a subscription.

For example, by signing up for a Gmail account through Google, you're accessing the application via the SaaS model. Yet when a larger corporation such as a university contracts with Google to run thousands of email accounts e.g. student email accounts, the solution changes from SaaS to cloud computing. Google is then managing the infrastructure headaches of running the IT infrastructure.

Services must be on a large scale to qualify as cloud computing. It is a matter of establishing how well the solution meets the needs of the user when determining whether SaaS or cloud computing is more appropriate.

### **1.7.3 Grid Computing**

Cloud computing is often confused with grid computing, a form of distributed computing whereby a 'super and virtual computer' is composed of a cluster of networked, loosely-coupled computers, acting in concert to perform very large tasks.

Grid computing has been used in environments where users make few but large allocation requests. For example, a lab may have a 1000 node cluster and users make allocations for all 1000, or 500, or 200, etc. So only a few of these allocations can be serviced at a time and others need to be scheduled for when resources are released. This results in sophisticated batch job scheduling algorithms of parallel computations.

### **1.7.4 Utility Computing**

Many cloud computing offerings have adopted the utility computing model. Utility computing can be analogized to traditional computing such as the consummation of electricity or gas. IT computing becomes a utility, similar to power, and organizations are billed either based on usage of resources or through subscription.

## **1.8 *Virtualization Processes***

Some of the service management processes involved with Virtualization include:

- Demand Management;
- Capacity Management;
- Financial Management;
- Availability Management;
- Information Security Management;
- IT Service Continuity Management;
- Release and Deployment Management;
- Service Asset & Configuration Management;
- Knowledge Management;
- Incident Management;
- Problem Management;
- Change Management; and
- Service Desk Function.

Each of these processes will be discussed further in chapter 5.

## **1.9 Introduction Review Questions**

1.29 Which of the following is not a main objective of virtualization?

- a) Increased use of hardware resources.
- b) Improved business flexibility.
- c) Reduced management and resource costs.
- d) Lowering IT management costs.

2.29 If a server fails due to a hardware fault, what happens to the virtual machines stored on it?

- a) They are always accessible.
- b) They can be migrated to another server.
- c) They are not able to be recovered.
- d) None of the above.

3.29 True or False?

Virtual machines are not separate entities from one another.

T      F

4.29 Which of the following is a true statement in relation to SaaS?

- a) SaaS is software that is delivered by many providers, yet managed by a sole provider.
- b) SaaS customers access computing resources on-site, whereas cloud computing is offsite.
- c) SaaS allows sharing of application processing and storage resources.
- d) SaaS does not allow the sharing of resources in a one-to-many environment.

*Answers to all questions can be found in Chapter 7.*

## 2 Common Terminology

Critical to our ability to participate with virtualization is the need to be able to speak a common language with other IT staff, customers, end-users, suppliers and other involved stakeholders. Below is a list of some of the most common terms that are related to virtualization.

<i>Term</i>	<i>Definition</i>
<b>Bare Metal Environment</b>	This describes virtualization software being installed directly onto physical hardware. It is installed in place of an OS and is the opposite of a hosted environment.
<b>Capacity Planning</b>	This involves working out how many physical resources a machine will need in order to host a set number of virtual machines. It is an important process that is carried out at the beginning of a virtualization project.
<b>Cloud Application</b>	Leverages the cloud in software architecture and often eliminates the need to install and run the application on the customers own computer, thus alleviating the burden of software maintenance, ongoing operation and support.
<b>Cloud Architecture</b>	The systems architecture of the software systems involved in the delivery of cloud computing.
<b>Cloud Client</b>	Computer hardware and/or computer software which relies on the cloud for application delivery, or which is specifically designed for delivery of cloud services.
<b>Cloud Computing Provider</b>	Owns and operates live cloud computing systems to deliver service to third parties.
<b>Cloud Infrastructure</b>	The delivery of computer infrastructure as a service.

<b>Cloud Platform</b>	Facilitates the deployment of applications without the cost and complexity of buying and managing the underlying hardware and software layers.
<b>Cloud Service</b>	Software systems designed to support interoperable machine-to-machine interaction over a network which may be accessed by other cloud computing components, software or end users directly.
<b>Cloud Storage</b>	The delivery of storage data as a service and often billed on a utility computing basis.
<b>Emulation</b>	This involves using software to run operating systems or applications on hardware that they were not originally intended to run on. For example, virtual machine software emulates x86 hardware to allow operating systems to be installed on them, even though they are not physical machines.
<b>Grid Computing</b>	This is a relatively new model based around virtualization and it involves using the processing power of a large number of geographically separate but networked machines, to complete complex tasks in parallel. Grid computing is often used by scientific institutes to use the idle processing power of user's machines. Some of the institutes that use grid computing include SETI and more recently the PlayStation 3 for scientific research.
<b>Guest OS</b>	This is the operating system that is installed within a virtual machine or a partition. When using virtualization the guest OS can differ from the host OS. However when using partitioning, the guest OS must be the same as the host OS.
<b>IaaS (Infrastructure-as-a-Service)</b>	The delivery of computer infrastructure as a service.
<b>Hardware-Assisted Virtualization</b>	This is similar to full virtualization, the only difference being that this type of virtualization is assisted by hardware (most often the processor) and is not solely software based.

<b>Host OS</b>	This is the operating system that is installed on a physical machine, which is hosting one or more virtual machines.
<b>Hosted Environment</b>	This describes a situation where virtual machine software is installed on top of a host operating system. It is the opposite of a Bare Metal Environment.
<b>Live Backup</b>	This involves copying an entire virtual machine, when logged in at the host OS level, for backup purposes.
<b>Live Migration</b>	This involves copying an entire virtual machine, when logged in at the host OS level, to another machine at the host OS level. This is often done when a physical server needs to be upgraded or taken down for maintenance.
<b>P2V (Physical to Virtual) Migration</b>	This describes the process of moving an entire operating system along with all its applications, from a physical machine to a virtual machine (or partition). This type of migration is carried out without having to reinstall anything, therefore it is quite efficient. P2V is the opposite of V2P (Virtual to Physical) migration.
<b>PaaS (Platform-as-a-Service)</b>	An entire virtualized platform that includes one or more services, operating systems and specific applications.
<b>Partial Virtualization</b>	This is a type of virtualization that only partially simulates the physical hardware of a machine. Because of this fact, most partially virtualized machines cannot run entire operating systems. However partial virtualization is easier to implement than full virtualization and many applications can still run on a partially virtualized machine.
<b>Partitioning</b>	This involves creating multiple 'virtual hard drives' from one physical hard drive, called partitions. Each partition can have the same host OS installed but unlike virtual machines, they



are not independent from one another. This means that partitions content can be accessed by a user who is logged in to the host OS of a completely separate partition.

**Partition**

This is a virtual hard drive that has been created from a physical hard drive, which can have a host OS installed on to it.

**RAID (Redundant Array of Independent Disks)**

RAID is a form of virtualization because it uses multiple disks, to simulate one large disk.

**SaaS (Software-as-a-Service)**

The ability to access software over the Internet as a service.

**V2P (Virtual to Physical) Migration**

This is the exact opposite of a P2V migration. It involves moving an entire operating system along with all of its applications from a virtual machine to one or more physical machines, without having to reinstall anything.

**V2V (Virtual to Virtual) Migration**

Similar to a P2V and V2P migration, it involves copying a virtual machines OS and applications but to another virtual machine. Again the process is completed without having to reinstall anything.

**Virtual Machine**

This is an isolated virtual environment within a host OS. All of a virtual machine's hardware such as processor, memory, disks, RAM, etc, are emulated and managed through the use of a virtualization application. This virtualization application is installed on the host OS, while guest OS' are installed within virtual machines.

**Virtual Machine Monitor (VMM)**

This is a simulated layer in between a guest OS and a machines physical hardware. It deals with hardware requests from a guest OS on a virtual machine, and responds to them in the same way that the physical hardware actually would. A VMM is also some times referred to as a Hypervisor.

**Virtual Memory**

This describes the process in which data is made to appear as though it is in contiguous order, even though it may be scattered in many locations across a machines RAM and Hard Drive.

**Virtualization (Full Virtualization)**

This describes the process which allows an operating system to be installed in an isolated virtual environment (i.e. a virtual machine). All the hardware resources of a virtual machine are emulated, while the physical resources are shared among one or more virtual machines.

## ***3 Virtualization Technologies***

VMware and Microsoft have a strong footing in virtualized technology and are perhaps the most well known companies involved with virtualization. However, now that virtualization technology has begun to stabilize and its advantages have been proven, plenty of other companies have started emerging. Below is a comprehensive guide on some of the most popular and widely used virtualization technologies as of 2008.

### ***3.1 Ubuntu (Server Edition)***

Ubuntu is a popular community-developed, Linux-based operating system. There are various editions of Ubuntu, such as Desktop Edition, Server Edition, etc. The great thing about Ubuntu is that it is open source and completely free. It comes with a number of pre-installed applications and a GUI (Graphical User Interface), which is one of the many reasons why it is so popular.

Ubuntu (Server Edition) also offers Kernel-based virtual machine (KVM). KVM is basically Linux's own full virtualization solution (for example Microsoft Virtual PC is Microsoft's virtualization solution). Due to the open source nature of Linux, many developers have optimized and continuously updated Linux's KVM technology. KVM can work with both Intel's (Intel VT) and AMD's (AMD-V) hardware virtualization technologies. Hardware virtualization is simply a virtualized software environment, which has received some help in virtualization, from hardware (usually the processor). This makes the virtualization process much more efficient. Most modern x86 processors for both Intel and AMD will have virtualization technology built-in.

KVM is a more than capable virtualization technology and can run multiple virtual machines with a wide range of guest operating systems installed. Also because KVM is implemented within the operating systems kernel, it is much more efficient and requires fewer resources than an application-based virtualization solution would. The kernel is a vital part of an operating system, which provides communication between the hardware and software components of a machine. KVM requires far fewer overheads than other application-level virtualization technologies, because it is a part of its operating systems kernel and can communicate directly with the hardware that it is running on.

KVM also uses a common interface known as VirtIO, for the transfer of data between virtualization technologies and its physical disk or network. This means that far less emulation needs to take place, for simulated hardware. For example, traditionally a virtual NIC (Network Interface Card), along with its operating system's entire kernel would have to be simulated, so that it could transfer data between itself and its physical NIC. However using VirtIO, an entire virtual kernel does not have to be simulated, because VirtIO provides a common interface for the transfer of data between many different operating systems.

### **3.2 Altiris**

Altiris is an organization that was recently acquired by Symantec (the creators of Norton Anti-Virus) in 2007. They specialize in management software that makes it easier for administrators to manage IT environments. They also offer virtualization software known as Altiris SVS (Software Virtualization Solution), best of all this application is free for private use.

SVS allows you to install software applications on to your machine but in a virtualized area. This means that applications are not actually installed on to your physical machine, instead they reside in a virtual location. This virtual location can then be activated or deactivated at will.

There are many benefits to this. For example unverified software could be tested in this virtual location. Also, software conflicts could be avoided and many different versions of the same program could be installed onto a single machine.

Using SVS is simple enough. The first thing that you have to do is create a virtual layer. A virtual layer is simply the virtual space which applications will be installed on. SVS allows you to create multiple virtual layers on a single machine. Once a virtual layer has been created, all you need to do is browse to the installation file of a particular program and select it from within the virtual layer that you just created. You will now be able to activate or deactivate this application.

When you activate an application, it will appear to your machine as though it has been installed normally. For example the application will show up in your start menu, any required registry changes will be made, etc. However once you deactivate the application, any changes that were made when the application was activated, will be gone. It will

appear to your physical machine as though the application was never even run on your machine.

This is great because no changes are actually made to your host operating systems environment. It can even help to make a physical machine more secure. For example if an application that contained spyware was accidentally installed on a machine, it would normally take a lengthy and careful scan to remove this spyware. However using SVC, all an administrator has to do is deactivate the application and then get rid of it from within SVC.

### **3.3 Windows Server**

Windows Server is Microsoft's premier server OS. Its latest version, Windows Server 2008 is Windows first OS that is integrated with full virtualization. This feature is known as WSV (Windows Server Virtualization).

Along with its previous server roles, Windows Server 2008 has a new Hyper-V role. Selecting this role allows you to create a virtualized server environment. Within this environment you can then create multiple virtual servers. Hyper-V is much more efficient than stand alone virtual server applications because it is actually integrated into the Windows Server 2008 operating system.

WSV offers many of the benefits that you would expect to receive from virtualization technology, such as rapid server deployment and the ability to take snapshots for backup purposes, etc. However it also offers a few extra features. For example it can offer up to 32 GB of RAM and up to 4 processors per guest OS. It also provides support for virtual LANs and can run both 32bit and 64bit virtual machines. WSV has plenty of resources and can be flexibly administered by IT professionals.

WSVs main competitor is VMware ESX edition. They are both very strong virtualization technologies and both companies have a solid footing in the virtualization market area.

WSVs main advantage is that it supports more hardware than VMware ESX currently does. As long as you can install Windows Server 2008 on a machine, you should be able to run WSV.

VMware however, has a much larger feature set than WSV and it is a tried and tested virtualization solution. It will be interesting to see how things will develop between these

two companies in the future and whether VMware can continue to remain the market leader in virtualization technology.

### **3.4 Software Virtualization**

Software virtualization is the virtualization of applications or computer programs. One of the most widely used software virtualization programs is SVS (Software Virtualization Solution), developed by Altiris.

The concept is similar to hardware virtualization where physical machines are simulated as virtual machines. Software virtualization involves creating a virtual layer or virtual hard drive space where applications can be installed. From this virtual space, applications can then be run as though they have been installed onto host OS.

Once a user has finished using an application, they can 'switch it off'. When an application is switched off, any changes that the application made to the host OS, will be completely reversed. This means that registry entries and installation directories will have no trace of the application being installed or executed at all.

Software virtualization offers many benefits. Some of these benefits include:

- The ability to run applications without making permanent registry or library changes;
- The ability to run multiple versions of the same application;
- The ability to install applications that would otherwise conflict with each other (by using multiple virtual layers); and
- The ability to test new applications in an isolated environment.

Software virtualization provides many benefits and is easy to implement. A fantastic advantage is that you can try out software virtualization yourself by downloading Altiris's SVS application completely free.



### 3.5 VMware

VMware is one of the most widely known virtualization companies. Its brand is easily recognizable and they offer a number of virtualization programs. Their most popular virtualization applications are briefly detailed below.

#### **Desktop Editions**

**VMware Workstation** – Initially launched in 1999, VMware workstation is one of the longest running modern day virtualization applications. It allows users to create multiple x86-based virtual machines on a single physical machine. A wide number of guest operating systems such as Windows, Linux and MAC OS X can then be installed on to these virtual machines.

**VMware Fusion** – This is similar to VMware Workstation, the only difference is that VMware Fusion was designed for users of the MacIntel hardware platform. It is fully compatible with all virtual machines created by other VMware applications.

**VMware Player** – This application is a freeware application and is offered to users who do not have a licence to run VMware Workstation or VMware Fusion. Unlike the other two applications, VMware Player cannot create virtual machines, however it can run them.

#### **Server Editions**

**VMware ESX Server** – This is an enterprise level virtualization product that is offered by VMware. It does not require a host OS to be installed as it is installed directly onto a server's hardware (i.e. is a bare metal virtualization solution). This is unlike the desktop editions which are installed as applications from within their host OS. VMware ESX Server is much more efficient than other virtualization technologies because it has lower system overheads and interacts with its hardware directly.

**VMware ESXi** – This application is similar to the VMware ESX Server application. The only difference is that it takes up less memory, because its Service Console is replaced with a simpler interface. As of July 2008, VMware ESXi is available to download for free.

**VMware Server** – This is VMware's free server virtualization application. It is an application which needs to be installed onto a host OS that is either Windows or Linux based. Due to this fact it is not as efficient as the other VMware server editions, which are installed

directly on their hardware. However VMware Server does allow you to create multiple virtual machines which can have a number of different guest operating systems installed.

As you can see, VMware offer a number of different virtualization applications. Each one has its own advantages and disadvantages, depending on the scenarios they are used in. VMware Workstation would be best utilized in an environment which contains multiple end-user desktops. These desktops could be virtualized into a few physical machines. VMware ESX Server would be best used to create high performance virtual servers which provide important services. VMware Server is offered for free, however it is not as efficient as VMware ESX Server. However it would still be great to use for the virtualization of less mission-critical servers. Which virtualization solution a business should go for really depends upon their current scenario; for example how large they are, whether their servers are resource-intensive and what their future prospects are. Once a business determines these things they can go about choosing a suitable virtualization solution. Although VMware is the dominant brand in virtualization, other companies are now starting to catch up. Microsoft has released their new Hyper-V software, which is a direct competitor to VMware's own ESX Server edition. At the moment, tests show that performance wise, ESX Server Edition currently leads the way. However many virtualization companies are improving their technologies and you can now even download powerful open source virtualization technologies for free.

### **3.6 Intel Virtualization**

Intel virtualization (Intel VT) is Intel's hardware assisted virtualization technology. Traditionally virtualization would be carried out completely by software but this put a great strain on hardware resources. These days, most processors will have virtualization technology built in to them. This puts less strain on a physical system because the software based virtualized solution is assisted by its hardware.

Intel VT technology utilizes a machine's hardware-based resources such as its processor, chipset and BIOS (Basic Input Output System), to offload some of its software based virtualization workload to its hardware. This type of technology can provide near native performance ratios, because virtualization is being carried out directly by hardware, rather than just software alone.

The benefits of Intel VT are similar to the benefits of other virtualization solutions. For example it allows multiple virtual machines to be deployed on a single machine. It allows

the virtual machines to be isolated from one another and so on. The only difference is that Intel VT is more efficient than just using software-based virtualization on its own.

Another great thing about Intel VT is that it will work seamlessly with existing software-based virtualization technologies. IT administrators do not have to make any changes to their infrastructure or their virtual machines, Intel VT will work automatically. It also allows you to run 64-bit guest operating systems on a 32-bit physical machine, and provides more security because using hardware leads to much better isolation of virtual machines.

The best thing for organizations is that all of the above-mentioned benefits can be achieved without making any changes or upgrades to their software. All they have to do is ensure that the processor they are using has Intel VT enabled.

### **3.7 Red Hat Virtualization**

Red Hat is a Linux-based operating system. It is similar to Ubuntu, in that it is a Linux-based, open source operating system. The first version of it was released in 1994 and it reached the end of its life cycle in 2003 (version 9 was its last version). It was then replaced by Red Hat Enterprise Linux, with 5.2 the latest version released, as of May 2008.

As with most modern day operating systems, the latest version of Red Hat Enterprise Linux has strong support for virtualization. As of version 5, Red Hat Enterprise Linux has full integration of server and storage virtualization. It can support servers of any size and provide an unlimited number of guest spaces (virtualized areas). Finally, it can offer up to four virtualized areas for desktop clients by selecting the Red Hat Enterprise Linux operating systems, Multi-OS option.

Red Hat offers support for two main types of virtualized guests. They are Para-virtualized guests and full virtualized guests.

Para-Virtualization offers the highest performance ratio and can work with any processor. However, Para-Virtualization is limited to having Red Hat Enterprise Linux versions 4 and 5 as guest operating systems installed onto its virtual machines. It does not currently support any other operating systems.

Full virtualization offers support for many more operating systems. It can support Red Hat Enterprise Linux Versions 3 to 5, as well as other third party operating systems. The only

drawback to Red Hat's full virtualization solution is that it requires hardware assistance. This means that it requires a processor with Intel VT or AMD-V technology. However, most modern day processors come equipped with hardware-based virtualization technology, so this is not too much of a problem.

Some of the main features that Red Hat Enterprise Linux offers include:

- The ability for live migration of virtualized machines;
- The ability to easily backup and restore virtualized machines;
- The capabilities to run 32-bit Para-Virtualized machines on x86-based hardware running in 64-bit mode; and
- The ability to easily manage virtual machines and guest OS through the use of graphical tools.

The benefits that Red Hat Enterprise Linux offers are similar to the benefits that other virtualization solutions provide. For example, it can reduce the number physical machines that a business has. Older operating systems and applications can be installed onto new hardware. Resources are more easily managed and uptime can be greatly increased. Finally, applications and data can be isolated within their own virtual spaces. This improves security and makes it easier for administrators to troubleshoot and isolate problems.

### **3.8 Softgrid Application**

Softgrid is Microsoft's own software virtualization program. It was acquired by Microsoft in 2006 and is now referred to as Microsoft Application Virtualization or App-V. It is similar in functionality to Symantec's Altiris Software Virtualization Solution (SVS).

Microsoft Virtualization Solution works by creating a virtual space on which software applications reside. This virtual space is referred to as a Softgrid SystemGuard. A SystemGuard virtualizes the services and elements that software applications usually use. For example it virtualizes certain System files, Registry entries, .ini files, DLLs (Dynamic-Link Library), system services (for example the cut and paste features used by applications) and lastly a user's specific profile data. Through the virtualization of all these application specific features, no changes are permanently made to a user's physical system or host OS.

Microsoft's App-V is made up of three major components. These components are the Softgrid management console, the Softgrid client and lastly the Softgrid sequencer.

## **Softgrid Management Console**

Anyone who has ever used Microsoft's Management Console (MMC) should feel right at home when using the Softgrid management console. From this console, you can manage and control the applications that you will use virtually. For example you can change an application's security setting (or the security group that they belong to), update your applications, or view individual statistics about them. The Softgrid Management Console is effectively an administrator's control centre, from which they can then deploy and install applications.

## **Softgrid Client**

The Softgrid Client is installed onto the end-user workstations. Administrators can control how much access end-users can have to the Softgrid Client, making it both a secure and robust solution. As well as its graphical user interface, the Softgrid Client also has a command line interface (CLI). The CLI can not only be used locally to configure the Softgrid Client but it can also be used to configure the Softgrid Client remotely.

## **Softgrid Sequencer**

The Softgrid sequencer basically allows you to encapsulate the installation data of an application into Softgrid's native installation format (a Softgrid package). Once a Softgrid package has been created, it can be deployed and installed onto multiple clients without any additional input.

Microsoft's App-V solution provides many benefits.

Firstly, applications can quickly and easily be deployed onto multiple workstations at once. This saves a lot of time for administrators, because they no longer have to carry out repetitive installations.

The updating of applications is also made easier because now, all an administrator has to do to upgrade an application is upgrade it once in the Softgrid Management Console. All users who have access rights to that application will now see the updated application installed on their workstations.

Sometimes old or even malicious applications need to be removed entirely from a network of computers. Without the use of Microsoft's Softgrid Management Console, this would be a long and tedious task but a few clicks using the management console, and the required application is completely gone without a trace.

Software virtualization provides many benefits and is a technology to definitely keep an eye out for in the future. Just like Server/Workstation virtualization, it is a technology that can completely change a businesses infrastructure for the better.

### **3.9 Linux Virtualization**

Linux is a Unix-like operating system and is one of the most well known freely distributed open source collection of software programs. Due to its open source nature, many different types of Linux operating systems exist.

Some of the most well known Linux OS distributors include Debian, Red Hat and Ubuntu.

Most Linux-based operating systems will either have some form of Para-virtualization or hardware assisted virtualization technologies. Linux operating systems also offer a few different virtualization type technologies, when compared to other operating systems. These virtualization technologies will briefly be described below.

#### **Para-virtualization**

Traditionally, x86 hardware was not able to adopt virtualization-based strategies. This is because a CPU processes instructions differently based on whether it is in privileged mode or not, and a virtualized processor can not run in privileged mode (as this could pose a security risk by giving it more priority than the physical processor).

Due to the above-mentioned reasons, virtualization software has to deal with the instructions from its guest machine and then replace them with another set of instructions to ensure that they can be correctly read by their machines physical hardware. Constantly rewriting the instructions from the guest machine's virtual processor creates a large overhead, which greatly slows things down.

Para-virtualization is simply the modification of an operating system (usually at the kernel level) so that it can more closely resemble its physical machines actual hardware, making it more efficient at virtualization.

Although a Para-virtualization solution is quite efficient, it can only support a few operating systems, which is its main disadvantage.

## **Hardware assisted virtualization**

This is the second type of virtualization solution that most Linux-based operating systems will use and it requires either an Intel VT enabled processor or an AMD-V enabled processor. Using these processors, instructions can be read from guest machine virtual processors, without any swapping of instructions.

This is obviously much more efficient than having to swap processor instructions and unmodified operating systems can be installed on to virtual machines, because the physical processor will deal with the virtual processors instructions (rather than the operating systems kernel).

There are currently two Linux-based technologies that can both do hardware-assisted full virtualization. They are Xen and KVM (Kernel-Virtual Machine).

## **Coopvirt**

Coopvirt (Cooperative Virtualization) is an interesting concept. It is a hybrid virtualization solution, using both Para-virtualization and hardware-assisted full virtualization. It incorporates hardware-assisted virtualization from the processor, while also using a tidy and efficient operating system. Coopvirt is still in its infancy, however it could greatly improve the efficiency of virtualization like nothing that has been seen before.

## **Containers**

This is sometimes referred to as operating-system level virtualization and is not a virtualization solution in its truest sense. Containers are separate user environments but they are all run under a single kernel. This means that they are not totally isolated from each other (not in the way virtual machines are anyway). However, the good thing about containers is that they have lower overheads in comparison to true virtualization solutions, because they do not need to emulate entire kernels, processors, etc. Different Linux distributions can also be installed into to different containers, which is another good thing about them.

As you can see there are a number of different Linux-based technologies. True virtualization can not be beaten when virtualizing hardware assets. However, containers also have their place in the market and could be used by smaller organizations or home users to create different workspaces within their physical machine, without having to use resource-intensive virtualization solutions.

### **3.10 Desktop Virtualization**

Desktop virtualization is the virtualization of end-user desktops. This makes it easier for administrators to manage end-user machines and to customize individual user experiences remotely. There are two kinds of desktop virtualization technologies in use today. One is Client-Hosted Desktop Virtualization and the other is usually referred to as Virtual Desktop Infrastructure (VDI).

#### **Client-Hosted Desktop Virtualization**

This type of virtualization involves installing virtualization software (for example Microsoft Virtual PC), on to the end-user's physical machine. On this machine, multiple virtual machines can now be created for different purposes. For example, one virtual machine may be used as a test environment for new software, while another may only be accessible by certain users and restricted for others.

The advantages of this kind of virtualization are obvious; for example multiple desktop environments can be run on a single physical machine all isolated from one another. However, one disadvantage of using Client-Hosted Desktop Virtualization is that all virtual machine files are visible to any user who is using the physical machine on which the virtual machines reside. In a security-critical environment, this could pose a security risk because anyone logged into the host OS could effectively copy entire virtual machines, to a portable hard drive. The solution to this problem is using VDI.



## **Virtual Desktop Infrastructure (VDI)**

VDI is a relatively new type of virtualization technology, however, it is based around a traditional thin client-server model. It involves storing end-user desktop environments on virtualized machines, which are stored on servers. A user can then access their desktop environment through the use of end-user PC's or thin client machines.

VDI is similar to terminal service remote desktop access technology, its only difference being that it involves a user accessing a fully virtualized end user environment.

There are many benefits to using VDI. Administrators can centrally manage and control desktop environments. Security is improved because all virtual machines reside on centralized servers. The only real disadvantage is that VDI requires powerful server hardware, due to the extra load that is placed on a server, which has to carry out both virtualization and manage incoming connections to these virtualized desktop environments.

This also means that cheaper end-user desktop machines can be used because they require far less processing power. This means that a large business can actually save money in the long run by investing slightly more in server hardware but much less in end-user desktops.

VDI is best used in a business environment which has contracted or offshore workers. VDI offers all the benefits of traditional virtualization solutions but also offers the benefits of remote desktop access, and centralized desktop management for administrators.

### ***3.11 Hardware Virtualization***

Traditionally, full software-based virtualization was difficult to implement efficiently, due to the way x86 hardware and processors operate. It was possible, but there were massive overheads when a virtualized machine would communicate with its physical processor. Two solutions were developed to help overcome these problems and make virtualization more efficient. These two solutions were Para-virtualization and Hardware virtualization. It should be noted however, that hardware-assisted virtualization is not an entirely new technology and was used by powerful mainframe computers decades ago.

As of 2005, both Intel and AMD have adopted hardware-assisted virtualization technology into their x86 processors so that are more virtualization friendly. Intel's hardware virtualization technology is known as Intel VT, while AMD's is known as AMD-V.

These processors work by carrying out some of tasks that would normally be performed by software known as Virtual Machine Managers (VMM, also sometimes called a Hypervisor). Because a processor is directly carrying out important virtualization processes (which would otherwise be carried out by software and then be passed on to hardware), there are far fewer overheads and the whole virtualization process is much more efficient.

Again there are many benefits to using hardware-assisted virtualization. One of the major benefits is that the amount of hardware emulation required is greatly reduced, thus improving performance.

Traditionally, x86 operating systems need to have direct access to hardware resources, so that they can run properly. In the past, software-based virtualization solutions would need to emulate hardware in its entirety, so that virtualization stability could be ensured. This wasted a lot of physical hardware resources which could be used elsewhere. Hardware-assisted virtualization greatly reduces the amount of hardware emulation that is required, because most of the time, the processor can interact directly with its virtualized environments. Obviously this leads to much more efficient virtualization.

The great thing about hardware-assisted virtualization is that it will work automatically, straight out of the box. You do not need to update your virtualization software or modify your operating system in any way. All you need to do is use a processor which supports hardware virtualization technology. This processor will then assist with virtualization automatically, as long as its virtualization technology is enabled, which is usually via a physical machine's BIOS.

### **3.12 Resource Virtualization**

Resource virtualization involves the virtualization of one or more IT-related system resources. It could involve the virtualization of specific system resources such as storage or network resources, or it could involve the virtualization of entire physical resources, such as servers or end-user workstations.

Organizations are constantly looking for ways to reduce the costs of their IT infrastructure. Resource virtualization is one of these ways. By grouping together and sharing IT resources, a business can greatly reduce their infrastructure and hardware costs.

Additional benefits are also gained through resource virtualization. For example a business's flexibility is increased because they can now control processing power, move resources between different servers with ease and even completely power-down machines which are no longer in use. Different applications and resources that would usually conflict with each other can now also be placed onto the same physical machine. Finally, resource virtualization makes it easier for administrators to manage and control their IT resources. For example administrators could manage and monitor multiple virtual servers from a single machine or they could manage and even clone a specific virtualized storage area with ease.

### **3.13 Processor Virtualization**

Processor virtualization simply involves the virtualization of a physical processor, so that multiple virtual machines can be run on a single physical machine.

Processor virtualization is not the same as multitasking or Hyper-Threading technology. Multitasking is a process in which multiple applications can be run by a single processor at the same time. This may sound like virtualization, but it is not.

The reason for this is because multitasking applications cannot access privileged areas of a system such as a systems input/output area. Only an operating systems kernel can access these areas of a system and in a multitasking solution, there is only a single kernel running. Virtualization involves using multiple kernels, meaning each virtual machine is completely isolated from the other, as if they are separate physical machines.

Hyper-Threading is also not the same as virtualization (although it has a similar concept to virtualization). Hyper-Threading involves making a single processor appear as multiple processors. It is a bit like hard drive partitioning, where a single hard drive is made to appear as separate physical hard drives. By doing this, a processor can process multiple threads (pieces of applications code) at the same time. However, Hyper-Threading is not truly a virtualization solution, because there is still only one kernel running between a Hyper-Threading processor and its software.

True virtualization offers a one-to-many relationship between hardware and software. That is, many operating systems, including their kernels (which means operating systems are completely isolated from one another), can run on a single machine.

These days both Intel and AMD have virtualization enhancement tools, built in to their physical processors. These tools are known as Intel VT and AMD-V, respectively.

This hardware virtualization technology (available since 2005) has made virtualization more efficient than ever before. The reason for this is because using a processors virtualization tools, an operating system has far less strain on it to emulate hardware correctly and it can use its resources elsewhere, such as on a users processes.

As processing power greatly improves and the cost of processors comes down each year, more and more people will eventually have machines which are capable of providing efficient virtualization solutions. This could eventually lead to a massive adoption of virtualization technology worldwide, not just by businesses but by home users as well.

### ***3.14 Application Virtualization***

Simply put, application virtualization is the encapsulation of applications from the operating system that they are installed on. A lot of the times, certain applications that share the same system resources will conflict with one another. For example, certain applications may use the same drivers or edit the same registry entries. Usually when this happens, one or both of the applications will malfunction and then crash.

One way for organizations to avoid the above mentioned problem, is through extensive testing. However, this is often time-consuming, can be quite expensive and also a waste of an organization's resources. If two applications do conflict with each other, then administrators will usually have to place them on separate machines or on separate hard drive partitions. Although this will work, it is not the most practical solution. For example, a user would have to move between two different machines or log on and off between different hard drive partitions to access the separate applications.

A more robust solution is using application virtualization. Application virtualization works by making separate copies of shared resources, for separate applications. For example each virtual application will have a copy of its own drivers, registry entries, DLLs and any other resources that would usually be shared with other applications. When a virtual application is run, it will call upon its own specific resources. Therefore conflicts are avoided and

applications can run in complete isolation from one another, even though they are installed on the same operating system. The isolated environments that these applications run in are often referred to as virtual bubbles.

Application virtualization offers organizations many benefits. Some of these main benefits include:

- The ability to run non-native operating system applications. For example, running Windows applications on a Mac through virtualization or vice versa;
- The ability to isolate malicious or damaging applications;
- The capability to run applications which would otherwise conflict with one another;
- The ability to quickly and easily deploy software applications onto a large number of machines; and
- The ability for administrators to centrally manage applications along with controlling user access rights and privileges to run them.

The only real disadvantage of application virtualization is that applications usually have to be packaged by virtualization software, so that they can be correctly executed and run in a virtualized environment. However, this is not usually a difficult process and the benefits of application virtualization far outweigh its disadvantages.

### ***3.15 Storage Virtualization***

Storage virtualization involves the virtualization of physical storage devices. It is a technique that allows many different users or applications to access storage, regardless of where that storage is located or what kind of storage device it is.

When storage is virtualized, it appears standardized and local to host machines, even though the storage may be distributed across many different locations and many different types of hard drives.

The great thing about storage virtualization is that it allows many different machines and servers to access distributed storage devices. However, a particular machine accessing a virtualized storage area will see one large storage area, as though it is a single massive hard drive, rather than a load of scattered hard drives.

Other benefits offered by storage virtualization include the ability for administrators to mask particular hard drives or storage volumes from particular machines. This obviously improves security and the ability to increase a storage volumes size in real time. Again,

this is very useful because if a server appears to be running out of space on their virtualized storage area, an administrator can increase its size immediately with just a few clicks.

One of the most widely deployed storage virtualization technologies is a SAN (Storage Area Network). Just as its name suggests, a SAN is a large network of storage devices. These storage devices which are usually held in a rack are independent of any servers or machines, and are instead directly connected to an organization's network.

Through the use of Storage Area Networks a business can improve their flexibility. For example, a Storage Area Network's size can easily be increased by adding additional hard drives or storage volumes. They can be used as an efficient backup solution; for example by backing up data in a remote location away from business enterprises. Lastly, Storage Area Networks can provide better storage volume utilization. This means that instead of having one hard drive per server, a server can spread their data across multiple different hard drives or storage volumes. This is a much more efficient use of hard drives, because a single hard drive is not being constantly written and rewritten to. Instead, multiple hard drives share the load, which should increase the lifespan of individual hard drives.

### **3.16 Virtualization Density**

Virtualization density is a phrase that can be used with a number of virtualization technologies. It often refers to the number of virtual machines that can fit within a single physical machine or rack. The more virtual machines that a single physical machine can support, the more dense it is said to be.

It also sometimes refers to an organization's ability to run a large network infrastructure on fewer physical machines through the use of VLANs (virtual LANs) and full virtualization technologies. The term virtualization density may also be used to describe the number of virtualized applications that can run effectively on a single machine.

Generally, when the term virtualization density is used, it is used to describe how many of a particular virtualized technology (whether they are virtual machines, applications or some other quantifiable aspect) can run efficiently within their particular environment. The more quantifiable aspects that can fit within a particular environment, the denser that environment is said to be.

### **3.17 Para-Virtualization**

As mentioned previously, x86 processors are difficult to emulate because of the way in which they work. A host operating system's kernel needs direct access to its hardware in order for instructions to be interpreted correctly.

In a virtualized environment a guest operating system's kernel does not have direct access to its physical processor, therefore certain instructions have to be changed or replaced, so that there is a correct 'understanding' between the guest OS and its physical processor.

This modification of instructions was traditionally carried out by software. However, normal operating systems were never originally designed with virtualization in mind. Therefore they were not very efficient in a virtualized environment and produced massive overheads. This was one of the reasons why virtualization was slow and just not feasible for businesses until recently.

Two solutions were introduced which made virtualization much more efficient. These solutions were hardware-assisted virtualization and Para-virtualization. Hardware-assisted virtualization, as its name suggests, used physical hardware to take away the strain of virtualization from its software and operating system.

Para-virtualization on the other hand did not receive any additional hardware support. Instead it involved modifying an operating system's kernel to make it far more efficient.

An operating system's kernel would be modified so that it could run more smoothly in a virtualized environment. However, the disadvantage to this is that only certain operating systems can be used when Para-virtualization is implemented. Hardware-assisted virtualization is the way into the future as it requires no software modification whatsoever and will work with any operating system that supports x86 hardware.

### **3.18 OS Virtualization**

Although similar to full virtualization, OS-level virtualization is actually quite different. Full virtualization involves the virtualization of a machine's entire hardware. Each virtual environment is then run under its own operating system and more importantly its own kernel.

OS virtualization is different in that separate virtual environments with their own separate kernels are not created. Instead, OS virtualization works by running virtual environments (known as containers) under a single kernel.

Each container environment within an OS virtualization solution will be isolated from other containers and will look and act like a physical server. A container environment can then run applications and accept the workload of its physical machine. The end result of OS virtualization is effectively the same as full virtualization but as you can see, the process for each solution is actually different.

OS virtualization has a number of practical uses. In fact, it has been used in virtual hosting environments for many years. Virtual hosting involves hosting more than one domain name on the same physical machine. By using OS virtualization, web hosts can create secure isolated environments for different domain names. This is obviously advantageous because otherwise the resources of a single machine would be wasted if it could only be used as a host for one domain name.

Other benefits of using of OS virtualization include the separation of applications along with the ability to more easily manage resources. For example, using OS virtualization, you could separate or group applications into different containers. Software resources would also be more manageable because administrators would be dealing with smaller 'chunks' of resources, rather than entire groups of resources under a single environment.

OS virtualization sounds great and you may be wondering why most organizations today are using full virtualization solutions rather than OS virtualization solutions. Both solutions provide similar end results, however there are distinct differences between the two and OS virtualization does have its own set of pros and cons.

The major advantage of OS virtualization over full virtualization solutions is that it is far more efficient. OS virtualization has very little overheads because it does not need to emulate hardware. Communication between hardware and software is carried out by a container's host operating system's kernel, so again there is very little overhead. However, OS virtualization does have its disadvantages. Firstly, OS virtualization cannot run operating systems which are different from its original host operating system. If you want to run a Linux-based environment within a Windows operating system, then OS virtualization is no good to you.



Container environments also have a number of restrictions within them. For example, a container cannot modify its kernel directly, it cannot mount or dismount file systems and it cannot carry out other top level actions. A full virtualization solution on the other hand, gives a user a completely unrestricted environment on which many different operating systems can be installed.

In the end it was the flexibility that full virtualization solutions offered which made them the standard solution for virtualization. Along with hardware-assisted virtualization and Para-assisted virtualization technology, full virtualization is now just as efficient as OS virtualization. However OS virtualization is a technology that is still widely used; for example in web hosting environments and it will continue to be used in the future.

### **3.19 Virtualization Software**

Virtualization software is any kind of software that deals with the emulation of hardware and the splitting up of both physical and software-based resources.

Some of the most popular bare metal virtualization software applications include VMware ESX server and Windows Server Hyper-V. Bare metal virtualization applications are installed directly onto physical hardware in the same way that normal operating systems are installed. Guest operating systems will then be installed on top of a bare metal virtualization application.

The opposite of a bare metal virtualization application is a hosted virtualization application. A hosted application is installed directly onto a host operating system, just like any other application would be. This virtualization application can then run a number of guest operating systems under its host operating system.

Hosted virtualization is not as efficient as bare metal virtualization because a hosted application must communicate via its host operating system, before it can interact with its hardware. Some of the most popular hosted virtualization applications available today include VMware workstation and Microsoft Virtual PC.

Other virtualization software includes application virtualization software. Application virtualization software deals with the creation of separate virtual environments, from which applications can then be run within. These environments are completely isolated so that software conflicts can be eliminated. Another great thing about application virtualization is that when you install a program on your machine, there are no permanent changes.

Popular application virtualization software includes Symantec's Altiris SVS (Software Virtualization Solution), and Microsoft's Softgrid application.

Below is a list of some of the most popular virtualization software available today.

### **EasyVZ**

This is an open source application for Linux which offers a graphical user interface for users, helping them to manage their virtual private servers.

### **Fabric Server**

This is an application virtualization program. In other words, it removes an application from its operating systems environment and then places it into its own isolated environment.

### **FluidVM**

This is a virtualization management application. It offers support for a number of different virtualization strategies.

### **GridServer**

This is an application virtualization program created by DataSynapse. It virtualizes applications and isolates them from their hardware system and operating system resources.

### **HP Integrity Virtual Machines**

This is a hosted virtualization application, created by HP (Hewlett-Packard). It allows multiple virtual machines to run on any Itanium based server that has HP-UX installed as its operating system. An Itanium based server is one which has an Itanium processor installed. An Itanium processor is a 64-bit Intel microprocessor, which has a significantly different architecture than a standard x86 based microprocessor.

## **Hercules Emulator**

This is an open source emulator that can be run on Windows, Linux and MAC OS X-based machines, as well as on many other operating environments. It allows software that has been designed for mainframe IBM machines to run on cheap personal computers.

## **Hyper-V (Windows Server Virtualization)**

This is Microsoft's main virtualization software. It is available in two versions. It is available within Windows Server 2008 and is also available as a stand-alone bare metal solution.

## **HyperVM**

This is a Linux-based virtualization management program. It can support a number of different virtualization technologies across a number of different platforms and machines.

## **Leostorm P>V direct**

This is a physical-to-virtual conversion tool. It can convert physical Microsoft machines to virtualized VMware machines or Microsoft virtual servers.

## **Leostorm Virtual Desktop Connection Broker**

This is a virtualization application that can be used to monitor and control a number of virtualized machines, which fall under the same virtualization application. For example if you had a Microsoft Virtual PC solution, you could use the connection broker to manage all of the virtual machines that fall under this application.

## **LivePC**

This is a virtual PC which you can create and then share.

## **Mac-on-Linux**

This is an open source Linux-based application that can create virtual machines, which can run on PowerPC-based machines that have some form of Linux installed. It allows you to run Mac-based virtual machines within a Linux operating environment.

**Mac-on-Mac (MoM)**

This is an application that allows you to run numerous different Mac- and Linux-based operating systems within a separate Window on a Mac OS X-based machine.

**Microsoft Virtual PC**

This is Microsoft's hosted virtualization application. Only supported Windows operating systems can officially run within it. Although some Linux-based operating systems can run within a Microsoft Virtual PC, they are not officially supported and stability is not guaranteed.

**Microsoft Virtual Server**

This is similar to Microsoft Virtual PC, the only difference being that Microsoft Virtual Server was developed for use with server-based hardware. It also officially supports a wider range of operating systems, for example a number of Linux-based operating systems such as Red Hat Linux and Suse Linux.

**MojoPac**

MojoPac has a unique concept. It is a virtualization application that allows you to run your virtual PC on hard drives, USB flash drives, iPods and even mobile phones. It can run any application that a normal PC could, including high performance applications such as games. The great thing about it is that it offers the many benefits of virtualization, while also offering portability.

**MokaFive (Moka5) Live PC Engine**

This is a virtualization management application, which specifically deals with LivePC virtualization technology. Live PC Engine allows a user to make, run or share virtualized LivePCs.

## **Oracle VM**

This is the Oracle Corporations server virtualization software. It offers support for both Windows- and Linux-based operating systems and has its own built-in web browser management console.

## **Parallels Desktops for MAC**

This is a Mac-based full virtualization program which is produced by Parallels, Inc. It offers support for a number of 32-bit x86 operating systems such as Linux distributions, Windows, Solaris, etc. However it currently does not offer support for 64-bit operating systems.

## **Parallels Server for MAC**

This is basically the same as Parallels Desktop for Mac application, however it is intended for use with servers. It can support a number of operating systems, including both 32-bit and 64-bit operating systems.

## **Parallels Workstation**

Virtually the same as Parallels Desktops for Mac application, except that it is designed to run on Windows- or Linux-based operating environments. It can run multiple x86-based virtual machines at the same time. Again it offers support for a number of 32-bit operating systems. However it is unable to currently support 64-bit operating systems.

## **PearPC**

Pear PC is a PowerPC-based virtualization application which can run multiple operating systems such as Mac OS X, the UNIX based Darwin operating system and various different Linux based operating systems. PearPC can be run in both Windows- and Linux-based operating environments, along with a few other operating environments.

## **Q**

Q is based on the open source QEMU processor emulator application. It is freely available and runs within a Macintosh environment. It is used to run Windows or any other x86

based operating systems, on a Mac. Q is more user-friendly than QEMU because it has a graphical user interface for its management and configuration console, unlike QEMU which is solely command line based.

## **QEMU**

QEMU is primarily a processor emulator but it also offers tools that allow it to run a number of different guest operating systems. Therefore it effectively doubles up as being both a processor emulator and a hosted virtualization solution. Another great thing about QEMU is that it offers great portability. A virtualized QEMU environment can run on any PC, even on those machines where a user's rights may be restricted.

## **Quick Transit**

This is an application developed by Transitive Corporation which allows processor or platform specific software, to be run on another hardware or software platform. This is all done without having to recompile or make any changes to an application's source code. It is a Linux-based application but it has also been developed for other operating environments such as Apple, who refer to this technology as 'Rosetta'.

## **SIMH**

This is a portable virtualization application which can run on many different operating systems, including Windows, Linux, Mac OS X, OpenBSD and many more. It can also run on a number of different platforms, making it fairly versatile.

## **SVISTA (Serenity Virtual System)**

SVISTA is a virtualization application created by Serenity Systems International. It can run on a number of different x86-based physical machines and can create many different x86 virtual machines, which can run on a single physical computer. SVISTA virtual machines are capable of running Windows, Linux and BSD-type operating systems.

## **Sun VDI software**

Sun VDI (Virtual Desktop Infrastructure) is a desktop virtualization application that is created by Sun Microsystems. It provides desktop virtualization services by replacing multiple physical desktops, with virtual desktops that are stored on remote servers. Users can access their virtual desktop environment through the use of any thin client machine, as long as it has Java installed on it.

## **Sun xVM Ops Center**

This is Sun Microsystems' main management application. It can help administrators to manage both their businesses physical assets and their virtualized assets. It offers an administrator many benefits. For example, administrators can manage their organizations entire physical and virtualized assets from a single console. They can centrally manage updates and they can view detailed reports about specific or generalized machines.

## **Sun xVM Server**

Sun xVM Server, is Sun Microsystems' bare metal virtualization software and it can run on x86-64-based hardware systems. Sun Microsystems hope to add support for Windows and Linux distributions so that they can run as guest operating systems in the future. Sun xVM server does not have its own virtual hard drive format; instead it uses the same formats as VMware ESX Server and Microsoft Hyper-V so that these technologies can work in conjunction with one another.

## **Sun xVM VirtualBox**

Sun xVM VirtualBox is Sun Microsystems' hosted virtualization software, which runs on x86 hardware platforms. It can run multiple virtual machines within various different host operating systems. Some of these operating systems include Windows, Solaris, Linux and Mac OS X. It is similar to Microsoft's Virtual PC or VMware's Workstation solutions.

## **VDSmanager**

This is an application which allows administrators to manage Virtual Private Servers. It is a web-based control panel and provides administrators with a relatively easy-to-use graphical user interface (GUI). From here, administrators can then administer and create

Virtual Private Servers. Each VPS has its own set of emulated hardware and its own unique IP address.

## **VMmark**

This is a freely downloadable virtual machine benchmark application which was developed by VMware, Inc. It is used to measure the performance of a virtualized server when placed under a significant load. This is a useful tool for administrators, because it lets them see just how efficient their virtual machines really are.

## **VMware ESX Server**

This is a bare metal virtualization solution, which was developed by VMware, Inc. It is one of the most widely used virtualization solutions and can support a wide range of operating systems.

## **VMware Fusion**

This was developed by VMware, Inc, solely for Macintosh-based machines which use Intel processors. It is a hosted virtualization application, meaning it must be installed under a Mac-based operating system. Once this is done, Macintosh users can run various x86 based operating systems such as Windows, Linux and Solaris, within their Macintosh operating environment.

## **VMware Player**

VMware player is available to download for free by VMware, Inc. It is an application that can run virtual machines created by other VMware products. However, it cannot create the virtual machines themselves.

## **VMware Server**

This is a hosted virtualization solution which can create, run and edit virtual machines. It also offers a client-server-based service, meaning users can remotely access virtual machines using thin client technology.



## **VMware ThinApp**

Formerly known as Thinstall, this is an application virtualization program that virtualizes applications so that they become portable. Most Windows-based applications can be made portable using VMware ThinApp. However there are a few limitations.

VMware ThinApp cannot perform virtualization for any software program that requires a device driver to run. It cannot perform virtualization for any software that is copy protected via hardware or run any virtualized application which would not normally run under a particular operating environment. This basically means that you can't run Linux- or Mac-based virtualized application in a Windows operating system environment and vice versa.

## **VMware Workstation**

This is a hosted virtualization application, which allows users to create multiple x86-based virtual machines. VMware offers support for a number of different guest operating systems, including Windows, Linux and other BSD variations. One of its great features is that it allows users to take snapshots of their operating system while it is running. This means that a user's virtual machine can easily be restored or a person can easily return to a previous state, if it is necessary.

## **Virtual DOS Machine (VDM)**

VDM allows legacy 16-bit and MS DOS-based operating systems to be run on today's hardware and the latest line of Intel processors. VDM technology is currently available in all 32-bit Windows operating systems. However, it has been excluded from the latest line of Windows NT 64-bit operating systems, so they are unable to run 16-bit legacy DOS or Windows applications.

## **Virtual Iron**

Virtual Iron is a hosted virtualization solution. It can be installed directly onto supported hardware, without the need for the installation of another operating system first. Virtual Iron was one of the first companies to develop virtualization software which offered full support for hardware-assisted virtualization, using Intel-VT and AMD-V hardware-assisted virtualization technology. Because Virtual Iron utilizes hardware-assisted virtualization technology, it can run both 32-bit and 64-bit unmodified operating systems with close to native performance. Virtual Iron software is also supplied with a virtualization manager,

which can help administrators to control, change and monitor their virtualized environments.

## **Virtual Machine Manager**

This is a popular management application for managing virtual machines. It is provided for free with a number of Linux distributions, including Red Hat Enterprise Linux 5 and beyond, Fedora version 6 and beyond and Ubuntu 8.04 and onwards. It allows administrators to do things such as create and edit virtual machines, control virtual machines and view detailed statistics about performance and the utilization of resources by virtual machines.

## **VirtualBox**

This is a hosted virtualization application created by a German software company called Innotek. It can emulate standard x86-based hardware and run a number of different guest operating systems. Some of the host operating systems that VirtualBox can run on include Windows XP, Vista, Linux and Mac OS X, along with a few others. The major guest operating systems that it can support include Windows, Solaris, Linux and FreeBSD.

VirtualBox offers a number of benefits, for example it can run 64-bit guest operating systems, although this can only be done on 64-bit hardware. It also offers support for both Intel's and AMD's hardware-assisted virtualization technologies.

One disadvantage of using VirtualBox is that it uses its own proprietary format for its virtual hard drives. This means that VirtualBox virtual machines are incompatible with virtual machines created by other virtualization applications. However, VirtualBox can read and write to VMDK (Virtual Machine Disk Format) hard drives. VMDK hard drives are the virtual format VMware use for their virtual machine hard drive files. This means that VirtualBox can actually support virtual machines created by VMware, however VMware would not be able to support virtual machines created by VirtualBox, for now anyway.

## **Virtuozzo**

Virtuozzo is an operating system-level virtualization solution, created by Parallels, Inc. It has supported Linux operating systems since 2001 and Windows-based operating systems since 2005. Virtuozzo is able to create a number of virtual environments

(sometimes referred to as VEs or containers). These containers are completely isolated from one another and run as separate machines or servers on a single physical machine.

Because Virtuozzo is an operating-system-level virtualization solution, the operating systems within different containers, must match its physical machines guest OS. This means that if Virtuozzo is being used in a Windows XP operating system environment, then all of its containers can only run Windows XP operating system environments. However, Linux operating systems offer much more flexibility. Because Linux operating systems are all derived from a UNIX based operating system, they are effectively all the same (at least in underlying terms). This means that Virtuozzo running on a Linux-based host operating system can run a number of different Linux distributions for its container operating systems.

## **Vx32**

Vx32 is an application-level virtualization program. It can create isolated OS independent environments, in which x86-based software programs can run. It can run any plug-ins or extensions, written in any language as long as these applications are written to run on x86 hardware.

Vx32 has been stated as being similar to Java virtual machine. However using Vx32, users can run code written in any language, whether this code is safe or not. The main disadvantage for the Vx32 application is that at the moment, it is not very compatible with non-x86-based hardware.

## **Win4Lin**

Win4Lin is a virtualization solution restricted to working on Linux-based operating systems. It allows Linux users to run various versions of Microsoft Windows operating systems in a separate window, within a Linux-based operating environment. The main Windows-based operating systems that Win4Lin can support include:

- Microsoft Windows 95;
- Microsoft Windows 98;
- Microsoft Windows ME;
- Microsoft Windows 2000; and
- Microsoft Windows XP.

Win4Lin currently offers two main versions of its proprietary application: Win4Lin Pro Desktop and Win4Lin Virtual Desktop Server. Win4Lin Pro Desktop can be seen as being a tuned version of the QEMU virtualization application.

Win4Lin Virtual Desktop Server is a little different, although it is still based on the same underlying code. Virtual Desktop Server provides Microsoft Windows operating environments to users who are connected to a Linux server, using thin client technology.

## **XenMan**

This is a Linux-based virtualization management application. It provides users with a graphical user interface (GUI), which allows them to control the running of virtual machines, take snapshots and more.

## **Xenocode**

This is an application virtualization program which allows users to run software applications or programs virtually, without having to install them. Xenocode works by packaging applications into a single executable file. This file can then be executed like any other normal executable file, within a Windows operating environment.

There are no conflicts between Xenocode virtual applications and normal Windows applications which would normally share the same DLLs (Dynamic Link Libraries) or registry entries. The reason for this is because Xenocode virtual applications run in isolation from the rest of the system on which they are running, and they use a copy of their own resources such as DLLs so that they do not have to use their physical machines resources.

Xenocode offers a number of benefits to users. Its main benefits are:

- Applications do not have to be installed, therefore they can be rapidly deployed;
- Some older applications do not work within newer operating environments such as Windows Vista. Using Xenocode, users can run legacy applications within Windows Vista, which would not otherwise work;
- Sometimes a user may not have administrator rights on a particular machine. This would prevent them from installing software applications on to that machine and may also prevent them from running certain programs. Using Xenocode, a user can get applications to run, even on restricted machine, without having to make any changes to that machines host operating environment;

- Software conflicts are virtually eliminated because each Xenocode virtual application uses its own set of resources; and
- Security is also improved because no permanent changes are made to a host operating system when Xenocode applications are run. Also the applications run in their own isolated environment, meaning that any destructive applications can be tested and their effects isolated.

Xenocode is a well established application virtualization program. Although it was founded fairly recently (2002), it has been used by a wide range of high profile customers from a number of different areas. These clients include NASA, Siemens and Phillips, along with many more.

### ***3.20 Data Storage Virtualization***

As data storage requirements for businesses have steeply risen and continue to do so, managing this data has also become increasingly difficult. For example, backing up data can be quite a tricky task to do in a large organization. Over time, a number of technologies and solutions developed to overcome these problems and finally evolved into what we now know as 'Data Storage Virtualization'.

The first step taken towards developing a data storage virtualization solution involved eliminating the distinction between local and networked resources. By doing this, networked storage resources would appear as though they resided locally on a machine. These storage resources could also be accessed normally on individual machines in the same way that local storage would be accessed, making both local and networked storage devices appear as though they were logically identical (although to begin with, some configuration was required).

The advantages to this for businesses were obvious. Firstly a business's entire data storage pool could be held in a single centralized location. This would make data management for administrators a lot easier and would also help in copying, backing up or archiving data. Secondly, massive pools of storage area networks (SANs) could be created. This would allow organizations to have massive storage solutions which would just not have been practically possible in the past. For example, a server can host a number of hard drives, however, its physical storage space along with its hardware would limit how many hard drives it could support. Storage area networks on the other hand consist of a number of racks which contain many hard drives. There is no limit to how many racks an organization can have for a single storage area network. Really, the only

limiting factor for an organization is how much physical space they have or how much they are willing to spend on storage.

However, although storage area networks were great for organizations overall, they did not necessarily make things easier for end-users. This is mainly because a number of different protocols (set of rules) were used and administrators would still have to, at times, keep track of the physical location of data.

This is where data storage virtualization fully kicks in. Data storage virtualization solutions were introduced to standardize things, so that end-users did not have to configure anything or keep track of the physical processes of data storage. What this basically means is that a 'layer' was introduced between the physical and logical processes of data storage. This layer then standardized things, so that any physical storage solution could be automatically used in a logical solution, without any unnecessary configuration.

Data storage virtualization offers businesses a number of advantages. Some of these main advantages are briefly detailed below.

### **Provides easier management of data storage solutions**

These days, large data storage solutions are fairly complex. This makes it that much harder for administrators to manage an organizations data storage solution. By using Data Storage Virtualization technology, the complexities of dealing with massive data storage solutions for businesses are greatly reduced.

### **Allows storage technologies and solutions to be vendor independent**

In the past, using SANs that used hardware from different vendors or using different storage technologies would not be very feasible or easy to implement due the different set of protocols and rules that different vendors would use. However, storage virtualization makes it very easy to use different storage technologies with one another because it standardizes everything through virtualization.

### **Has the potential to save businesses a lot of money**

Physical storage management is still a tedious task for data centre managers. Their time and skills could be used in much more effective ways than having to do non-productive management tasks. Data storage virtualization will effectively eliminate the need for IT

professionals to manually deal with and manage physical storage, because this will all be taken care of through virtualization.

Another way that organizations can save money through data storage virtualization is that the hardware does not have to be vendor-specific. This means that they can save money by purchasing storage technology from different vendors. New technology purchased from a different vendor than the vendor that they purchased their existing technology from can also work seamlessly and with ease. This means that businesses do not have to change their entire storage infrastructure when upgrading their hardware, which again saves them a lot of money.

### **The main companies that have adopted data storage virtualization technologies**

The main companies involved in the development of data storage virtualization technologies include, HP, VMware, EMC and IBM. HP has been involved in the development of data storage virtualization solutions for businesses with enterprise network storage architectures.

VMware (a subsidiary of EMC) has also been strongly involved with data storage virtualization technologies. A number of their virtualization software solutions, such as VMware Workstation, have incorporated data storage virtualization in one form or another.

EMC are leading the way in this area of technology and their data storage virtualization solution is known as Information Lifecycle Management. EMC also have a hand in the development of data storage virtualization technologies for VMware.

Finally, IBM has developed a data server virtualization solution known as Total Storage. They have also developed a stand-alone storage area network (SAN) controller, which can perform various storage area network management tasks. This greatly reduces the workload for data centre managers, allowing them to use their time and resources much more productively.

### **3.21 Intel Virtualization Technology**

Intel virtualization technology specifically refers to Intel's hardware-assisted virtualization solution known as Intel VT. Hardware-assisted virtualization is effectively the same as other forms of virtualization, except that it is far more efficient. It is more efficient because

the physical processor takes a lot of strain of the host operating system's kernel and virtualization software.

Intel VT technology also reduces the amount of software-based hardware emulation that is required, because it carries out a lot of processes that would normally need emulating.

The benefits that Intel virtualization offers include:

- Reducing the number of physical machine that an organization has;
- Making management of a business's environment and infrastructure easier;
- Improving a business's availability and reliability while also reducing its downtime; and
- Improving resource utilization and reducing operating costs such as power, cooling and management costs.

There are many more but these are a select few.

Due to the fact that Intel VT takes a lot of the strain of software, it is capable of running and gives users the ability to run a very large number of different virtualized operating systems. It can even support operating systems which would not usually run effectively in a virtualized form without modification (a process known as Para-virtualization).

Intel VT has been designed to be easy-to-use and it will automatically work with virtualization software. Users do not have to do any configuring to get Intel VT to run with their existing software. All they need to do is enable Intel VT in their BIOS and the rest takes care of itself. Because of its ease of use, Intel VT has not only been provided in server microprocessors but also in desktop microprocessors. This means that even home users can take advantage of Intel virtualization technology. In fact, Intel VT is one of the many reasons why it has actually now become very practical for home users to run virtualization software on their machines, sometimes even at near native performance. Without Intel VT technology, this never would have been possible in the past, unless an operating system was specifically modified to run virtually (which is effectively what Par-virtualization is).

As hardware-assisted virtualization improves and Intel virtualization technology further develops, more and more home users can be expected to run at least some forms of virtualization software on their machines.



### **3.22 Thinstall Virtualization Suite**

Thinstall virtualization suite (now known as VMware ThinApp), is an application virtualization program, created by Thinstall Inc. In January 2008 VMware acquired Thinstall and the rights to the Thinstall virtualization suite, renaming it as VMware ThinApp. The main objective of VMware ThinApp involves the virtualization of normal applications or software so that they become portable and run on any machine or platform.

Just like other application virtualization programs, ThinApp allows users to run applications within a wide range of operating environments and platforms, without having to actually install the application in a traditional sense.

It works by packaging an application along with its own set of resources such as DLLs (Dynamic Link Libraries) and registry entries. This means that an application no longer has to use its host operating systems resources, allowing it to run within its own isolated environment, without having to install anything on to its host operating system's physical machine.

The virtual environments that applications run within are a combination of both physical and virtual resources. Because of this, applications can run normally because it appears as though they are actually running within their host operating system's environment, even though they are running in a completely separate and isolated container.

VMware's ThinApp application offers users many benefits. Some of these benefits are briefly detailed below.

#### **Allows users to run applications within non-native operating environments**

It allows applications to run in operating environments that they would not normally run in. For example users could run Windows applications within a Linux-based operating environment. This is possible because VMware ThinApp applications are packaged with all their OS native resources. This means that they can run normally within any operating environment.

#### **Contains virus outbreaks and the effects of malicious applications**

Because ThinApp applications are run from within their own virtualized container, they are completely isolated from the rest of their operating environment. This means that the

effects of any malicious or damaging software can be contained within its own environment or container, allowing the host OS to continue operating normally.

### **Allows conflicting applications to co-exist with one another**

Sometimes applications may use the same resources, for example the same registry entries or DLLs, in a Windows-based operating environment. When this happens, conflicts will usually occur and the applications will not be able to co-exist with each other or run on the same machine. By using ThinApp to isolate applications within their own environments and with their own set of resources, software conflicts are effectively eliminated.

### **Can work great on machines which may have older hardware**

Although a full virtualization solution offers many benefits, it is not always an entirely practical solution. For example it uses quite a lot of hardware resources, meaning that a full virtualized solution may not be able to run as efficiently on older hardware. This is especially the case with machines that have older processors, which do not offer hardware-assisted virtualization support.

Application virtualization provides a great middle ground. It offers many of the benefits of full virtualization while using far fewer hardware resources. This means that application virtualization solutions can run effectively, even on older hardware.

### **Easier operating system migration**

When having to migrate to a newer operating system, along with having to install the new operating system, most of the time many previous applications have to be reinstalled. This can often be lengthy and time-consuming process. Using Application virtualization software such as VMware's ThinApp, no applications have to be reinstalled. All an administrator has to do is copy these applications over again, which is much easier.

### **Rapid deployment of applications**

Again, administrators do not have to go through the time-consuming process of installing multiple applications on many machines. Instead all they have to do is copy the 'packaged' applications over, which can then be run as normally as though they have been installed onto that machine. Even more efficient is hosting these packaged applications on

centralized servers. These applications can then be accessed by users, via thin client technology.

As you can see, VMware ThinApp offers a user many benefits. Most Windows based applications can be virtualized, although there are a few exceptions.

- Any software applications that rely on device drivers can not yet currently be virtualized;
- Any applications that are bundled with their operating system (for example Windows Paint or Calculator), cannot be virtualized;
- Any software applications that use hardware based copy protection can also not be virtualized. This is because they depend on their physical hardware to actually run. Hardware-based copy protection is often used by the video games industry; and
- Software that would normally be unable to run within a particular operating environment would still be unable to run within that operating environment even if it was virtualized. What this means is that if you have an application that can run on Windows XP but not on Windows Vista (due to way that application accesses memory, for example), then even if that application is virtualized, it will still not be able to run within a Windows Vista-based operating environment. This is because even though the application is virtualized, it still accesses its machine's physical memory in the same way.

### ***3.23 Net Framework Virtualization***

Net framework virtualization (or Dot Net framework virtualization) involves running Dot Net applications without actually installing the dot net framework. The dot net framework is a library of pre-coded solutions to common software or programming problems. It is created by Microsoft and is used in conjunction with Microsoft operating systems.

Having the latest version of the dot net framework installed on your machine can provide you with many benefits. However, there are actually a number of problems and challenges that a user can also face and these problems are mainly due to software conflicts.

The first main problem that users may face when installing the newest version of the dot net framework is that some of their applications may stop working. Sometimes older applications may require an earlier version of the dot net framework in order for them to work correctly.

The second problem that users face stems from the first problem, in that a user may require different versions of the dot net framework so that all their applications can work. Running more than one version of the dot net framework can cause a number of problems and conflict with one another, reducing overall system stability.

Clearly there is a problem here then, due the fact that different applications often require different versions of the dot net framework to be installed, so that they can work.

Dot net framework virtualization solves the above problems in the following ways.

### **No need for installation**

Through the virtualization of the dot net framework (using some kind of application virtualization program), the dot net framework does not have to be installed on to a machine in the traditional sense. This means that no changes are made to a computer's host operating system environment and because the dot net framework does not have to be installed at all, it saves time.

### **Eliminating conflicts**

Because the dot network framework is not actually installed onto a machine and is run within its own isolated environment, conflicts are virtually eliminated. This also means that multiple versions of the dot net framework can be run on a single machine and therefore, applications that would not be able to co-exist with each other would be able to run without any problems.

### **Rapid and Easy deployment**

Application virtualization solutions usually package entire applications into a single executable file. This means no installation is required and administrators can copy over multiple versions of the dot net framework with ease. In fact, this is not even necessary because multiple versions of the virtualized dot net frameworks can be stored on a single server. Many users can then utilize these frameworks through the use of thin client technology.

Through the use of virtualization technology, Microsoft's dot net framework can be much more effectively utilized. Conflicts are eliminated and both new and old applications can

work on the same machine or within the same operating environment, without any problems.

### **3.24 *Windows Virtualization on Fedora***

Fedora is one of the many Linux-based distribution operating systems. In the past, running Fedora (or any other Linux-based distribution for that matter) and Windows on the same machine, would require the logical partitioning of that machine's hard drive.

This would create a dual boot situation, where one or the other operating system could be selected upon boot-up. However, both operating systems would not be able to run at the same time and you would have to reboot your machine before you could load up the other operating environment.

The solution to this problem was virtualization. By using virtualization, a Windows operating system could run within Fedora as a stand-alone application. This meant that both operating systems could now be run at the same time and there was no need to reboot or restart the physical machine when switching between operating systems.

Fedora uses KVM (Kernel based Virtual Machine) as part of its virtualization solution. This means that it fully supports Intel VT and AMD-V technology (Intel and AMDs hardware virtualization technologies, respectively).

### **3.25 Storage Virtualization Technologies**

Storage virtualization in its simplest form involves grouping together multiple physical storage devices (no matter where they are physically located), so that they appear as one big storage unit.

There are many different storage virtualization technologies including, NAS (Network-attached Storage), SAN (Storage Area Network) and iSCSI (Internet SCSI). These are described below.

#### **NAS (Network-attached Storage)**

A Network-attached storage device is effectively a server that is dedicated to providing file sharing services. It is not concerned with any other services such as e-mail or printer sharing; it only deals with file sharing.

NAS technology provides users with many benefits, for example it can allow additional hard drive space to be added to a network without having to shut down existing servers, thus improving the availability of data and increasing uptime.

When using NAS technology, storage does not have to be a physical part of a server. A server will still process all data requests that it receives but it is actually the NAS device that ultimately delivers data to a user.

NAS devices can be located anywhere within a LAN, meaning that a business can have the option of either centralizing or distributing their storage devices throughout their network. A single NAS can be made up of many NAS devices, meaning that a business can create massive pools of storage.

#### **SAN (Storage Area Network)**

A storage area network is effectively a sub-network that consists of only storage devices. Storage devices can consist of either a server which contains nothing but hard drives for storage, or a stand-alone rack of storage devices such as a rack of multiple hard drives.

SANs are designed in such a way that storage is accessible by all servers across an organization's LAN (Local Area Network) and their WAN (Wide Area Network). When additional storage is added to a SAN, it will immediately be accessible by servers across

an organization's entire network. Servers simply act as gateways between SANs and its end-users. They also help in the directing of traffic.

SANs provide organization with many benefits. One benefit is the ability to more efficiently use a server's resources. Because a server no longer has to deal with storing data itself, (it simply directs traffic where to go), a lot of its processing power is freed up.

Lastly, through the use of SAN technology, end-users have access to their organization's entire storage resources, rather than just the small amount of storage resources that would normally be accessible on a server (if SAN technology was not used). This basically means that end-users have much larger storage capacities than what normally be possible, without the use of SAN technology.

## **iSCSI**

iSCSI stands for Internet Small Computer System Interface. SCSI is basically an interface that is used for connecting peripheral devices (including hard drives) to a computer. It acts as a controller, meaning that it doesn't require software to handle data requests, as it deals with this itself.

The advantages of using SCSI rather than IDE (Integrated Drive Electronics), to attach a hard drive to a computer, is that SCSI offers faster transfer speeds. Especially in a server-based environment where multiple users may be trying to access a hard drive at the same time, SCSI definitely outperforms IDE.

SCSI also allows you to chain together multiple peripheral devices and connect them to a single SCSI interface. This means that using SCSI, you can connect many hard drives to a single machine, far more than what IDE would allow.

iSCSI expands on the standard SCSI technology described above by using IP (Internet Protocol) technology.

iSCSI can support existing Gigabit Ethernet technology and is able to connect to network devices such as switches and routers, as long as they provide an iSCSI interface. This means that iSCSI storage devices can be rapidly accessed across a network by separate users, as though they were being accessed locally.

As you can see, storage virtualization technologies offer many benefits to organizations. By making multiple physical storage devices appear as one or making distributed storage appear as though it is local, a business's storage solution can be a lot more flexible and manageable than it has ever been before.

### **3.26 Virtualization Level**

Virtualization levels refer to the different levels at which virtualization can be applied, when dealing with data storage virtualization solutions. These levels differ and can involve virtualization at the hard drive level all the way up to a full virtualized solution at the file system level.

There are four main levels which are described below.

#### **Server Level**

This level is fairly self explanatory. It involves the management of physical hard drives, within a server's operating system environment. For example it could involve logically partitioning a server's hard drive.

#### **Storage Device Level**

This is virtualization at the storage volume level and can involve techniques such as data striping or data mirroring. The most well-known storage device level virtualization technology is RAID. RAID 0 uses striping, which is a technique where data is written across multiple storage devices. However an end-user can access this data although it is stored on single local storage volume.

Other forms of RAID (such as RAID 1) use mirroring techniques where data is written on to two separate hard drives at the same time. Some forms of raid even use a combination of both striping and mirroring, making them an even more effective storage device level virtualization technology and backup solution.

#### **Fabric Level**

This level of virtualization involves allowing storage devices to be independent of physical hardware or servers. This means that massive pools of storage areas can be created, while logically appearing as a single storage solution. Usually a management interface is



required to manage storage systems at this level, because they are independent from server management. A storage area network (SAN) volume controller will carry out virtualization at this level.

## **File System level**

This level provides users with the most benefits because storage is completely virtualized. Data can be shared, protected or allocated from within software at this level, rather than at the physical hard drive level, which requires much more effort and resources. By using file system level virtualization the size of virtualized data pools can be manipulated at will, while complete clusters of data can be moved from one location to another, in the form of virtual hard drives, with complete ease.

### ***3.27 Security Monitoring and Virtualization***

Virtualization technology has made a lot of progress since it was first used conceptually in massive mainframe computers. There are more reasons today than ever before for businesses to migrate their physical infrastructure to a virtual one and a lot of companies have already made this move.

As businesses migrate from a physical to a virtualized environment, they can often forget to implement best practise security measures or feel that their old security measures are no longer needed. After all, virtualized environments do provide an extra layer of security when compared to just a physical infrastructure alone.

However, thinking like this is bound to leave some kind of security flaw and even the strongest systems can be brought down because of just one 'weak link'. It is important therefore, that a business puts just as much thought and effort into their virtualization security policies, as they would for their physical infrastructure security policies.

Detailed below are a few of the most common pitfalls or security issues that virtualized infrastructures are often faced with.

#### **All it takes is the infiltration of just one machine**

Firstly, when dealing with virtualized infrastructures, IT professionals should remember that all it takes is the infiltration of a single physical machine to bring down multiple virtual

machines. There is effectively a single point of failure when using a virtualized infrastructure.

### **The differences between virtual and physical are becoming less and less**

Virtual machines may often be seen as being 'artificial' and because of this administrators may overlook things that they wouldn't normally overlook, when dealing with physical machines. For example, administrators may not worry about installing anti-virus software on virtual machines or updating them with the latest patches, even though they would for normal physical machines.

This is just asking for trouble. In the past, administrators may have gotten away with such things because virtual machines were inferior to physical machines. However, as virtual machines begin to perform at near-native performance levels they are effectively becoming more and more similar to physical machines. This means that virtual machines are taking up the same traits as physical machines, including all of the flaws.

### **The failure of just one machine can cripple a business like never before**

In the past, if an important machine or server failed, it would often create quite a few problems for a business, especially if they didn't have a robust backup solution. However, other aspects of the business would still be able to continue operating normally. For example a business's print server may have failed but other services such as web access, file sharing, etc would continue working, as these services would normally be provided by separate physical servers.

With the advent of virtualization technology, separate services could co-exist with one another on a single machine. Of course this was great. It could save business's power, space, time and money.

However, it also created a single point of failure. If a single physical machine hosting multiple virtual servers fails, then the effects for a business can be damaging. In fact it can bring a business to a halt if no efficient backup solution is used.

As great as virtualization is, it has its own set of weaknesses and security flaws. However, every problem has a solution and listed below are some of the best ways to make a virtualized environment as secure as possible.

## **Keep all machines up-to-date**

Even when virtualization is not implemented, a business should ensure that all of their physical machines are updated with the latest operating system patches and application versions. However this should also be done for virtual machines and they too should be regularly updated. Using virtualization management tools, this process can be made a lot easier.

## **Keep applications to a minimum**

Again, this is something that is beneficial for both physical and virtual machines. Installing an excessive number of applications not only takes up space and resources but also makes it harder for administrators to troubleshoot if a problem does occur. You should only install applications that are necessary and any bundled or extra applications should not be installed.

## **Use firewalls on virtual machines**

A firewall is essential for any machine that is connected to the internet. Although administrators may think that virtual machines do not require firewalls, if these machines are connected to a network that is connected to the internet, then firewalls should definitely be installed. By using firewalls on virtual machines, they will be further isolated from one another, meaning that they cannot infect each other with viruses or other malicious software. Also, firewalls will shield these virtual machines from external machines outside of their own network, providing just as much security as a firewall would for normal physical machines.

## **Use Anti-Virus software on virtual machines**

Although virtual machines will isolate the outbreak of any virus infection from the rest of their physical system, they are still just as likely as a physical machine to actually receive a virus in the first place. Therefore anti-virus software should still be used to ensure that a virtual machine remains as virus-free as possible.

## **Don't browse the Internet from a physical host machine**

The internet is one of the most common sources for a virus infection. When using a physical machine that hosts many other virtual machines you should not browse the

internet. If a virus infects your physical machine's host operating system environment, then it has the potential to damage or even delete every single one of the virtual machines that are stored on that physical machine.

One great way to avoid any internet-related problems using virtualization is to create a single virtual machine that is dedicated for internet use. This can then be used alongside other virtual machines and if there is a problem on it due to a virus or malicious code, it can just be replaced by a completely fresh virtual machine.

### **Harden your physical machine's host operating system**

Server hardening involves making a server more secure by turning off or disabling any services that are not used. This should be carried out for both physical machines and virtual machines.

### **Power down both virtual and physical unused machines**

If virtual machines are not being used then they should be shut down or closed. The less virtual machines that are running, the less potential routes there are for a hacker to gain entry into a physical machine or network. Shutting down unused virtual machines also frees up memory and processing power which can be used for other virtual machines or processes.

### **Disable unused hardware**

If your virtualized environment will not be using USB technology, or for example optical drives, then these hardware features should be disabled from within your virtual machine. Again, this is similar to server hardening where a machine is made more secure by turning off unused features.

**Monitor both host machine and physical machine logs**

Logs keep track of computer transactions and important events. Most operating systems by default will note these events in a log file. Log files can be used by administrators to detect any suspicious activities or find security flaws that may need patching up. It is not uncommon for administrators to overlook the monitoring of their virtual machine's log files. This is because most administrators will already keep track of their physical machines log entries and feel it is unnecessary to also keep track of the log entries for each one of their virtual machines.

However, virtual machines are completely separate entities from their physical host machine. This means that any event taking place within a virtual machine that is logged will only be logged within that virtual machine's operating environment, and not by its host operating system or any other virtual machine. This is why administrators should monitor both the physical and virtual machine log files for each one of their virtual machines.

There are various ways to help secure and monitor virtualization solutions. Some of the most common ways were described above. As virtualization technology matures, some of its most common problems should be eliminated, whether through hardware or software means but generally, administrators should take the same precautions and security measures that they would normally take when securing physical machines. There is no need for organizations to change their security policies just because they are now operating virtually.

**3.28 Oracle Virtualization**

Oracle's main virtualization application is known as Oracle VM. This is a virtualization application that offers support for Oracle and even non-Oracle applications. Benchmark tests have shown that Oracle VM can be up to three times more efficient than other virtualization applications in certain aspects. It can support both Windows and Linux distributions as guest operating systems and also includes a web-based management console, making it easier for administrators to manage their virtual environment.

Along with their virtualization application, Oracle also offers what is known as Oracle VM templates. Oracle VM templates are effectively images (or snapshots) of pre-installed and pre-configured enterprise software. This makes setting up and configuring new machines much easier for administrators, because all they have to do is copy over an Oracle VM template which contains the software that they require.

Oracle's VM solution provides many benefits to administrators. It makes installation and deployment of software onto new machines much quicker and easier. It provides good efficiency and can support both Windows and Linux guest operating systems. Best of all, it is available to download for free.

### **3.29 Virtualization Technologies Review Questions**

1. Full software-based virtualization was hard to implement efficiently. What is one solution that was developed to overcome problems?
  - a) Desktop virtualization.
  - b) Processor virtualization.
  - c) Application virtualization.
  - d) Hardware virtualization.
2. What is not a benefit of Application virtualization?
  - a) Ability to separate or group applications into different containers.
  - b) Ability to isolate malicious or damaging applications.
  - c) Capability to run applications which would otherwise conflict with one another.
  - d) Ability to run non-native operating system applications.
3. True or False?

A host's operating system's kernel needs direct access to its hardware in order for instructions to be interpreted correctly.

T        F

4. Which of the following is not a virtualization level?
  - a) Server level.
  - b) File management level.
  - c) Fabric level.
  - d) Storage device level.

*Answers to all questions can be found in Chapter 7.*

## ***4 Accomplishing Virtualization***

### ***4.1 Migrating to a Virtualized Environment***

As virtualization technology has matured and its benefits have been proven, more businesses than ever before are now converting their physical infrastructure to a virtualized one. However, migrating from a physical to a virtualized environment is not always as straightforward as it seems.

The first thing that a business should do to ensure that they have a smooth transition from a physical environment to a virtualized one is determine what tools they will use to carry out their migration. During this initial planning stage, a business should also determine what they are actually going to be converting from physical to virtual.

The planning stage before a P2V (Physical to Virtual) conversion takes place, is a very important stage and can often be longer than the migration process itself. However, when a business makes a big change to their infrastructure such as this, good planning is essential to ensure that there are no unseen problems. Below are a number of things to consider and carry out during the initial planning stage before any migration takes place.

#### **4.1.1 Things to Do and Consider before Migration**

##### **What is going to be virtualized?**

Is a business's entire infrastructure going to be virtualized or just its servers? Is a business going to use application virtualization or storage virtualization technologies? These are important questions to which a clear answer must be given, before any migration takes place.

##### **Preparing for Downtime**

No matter how carefully a physical to virtual migration is planned, there is bound to be at least some downtime. During the planning stage a business should work out how long this downtime is and what they can do to minimize it or continue running with the very least amount of downtime.

Sometimes downtime can be as short as switching off a physical machine and replacing it with a new physical machine that hosts virtual machines. Other times it can be as long as waiting for the installation of a bare metal solution to be completed. No matter what the

scenario is, a business should take the necessary steps and precautions to prepare accordingly.

## **Hardware Requirements**

This involves working out whether virtual machines can run efficiently on existing hardware or not and if so, how many are required. This is important because if, for example, existing servers could only run two virtual machines efficiently, then the cost of migrating from a physical to virtual environment using existing hardware would be a waste because a ratio of 2:1 is not very good. In this case, the most efficient solution would be to replace existing hardware, so that five virtual machines could be run per physical server, for example. In the long run, this would be much more cost-effective than the first solution.

However, a business also wants to be careful that they do not overcompensate. For example, if a business only required five virtual servers per physical machine and they ordered in servers which were capable of running ten virtual machines, this would be a fifty percent waste of both money and resources.

## **Storage Requirements**

Virtual machines are encapsulated within files on a physical machine's hard drive. This means that hard drive or storage volume requirements are a big factor when dealing with virtualization. If storage is centralized, i.e. physically located within a server, then a specific storage volume or hard drive capacity would be required. For example, if a single physical server was going to be used to host five virtual machines each with a storage volume capacity of 100 GB, then at the very least, this server would require at least one 500 GB hard drive.

If storage was distributed, for example via the use of a SAN (Storage Area Network), then an organization would need to ensure that their SAN was large enough to accommodate all of its virtualized machines. If it was not large enough, then an organization could increase the capacity of their SAN by adding more storage volumes.

## **Ensuring that all Software will continue working in a Virtualized Environment**

There is nothing worse than finding out that important programs or software applications no longer work after a business's infrastructure has been completely virtualized. Certain software applications won't work in a virtualized environment and may not support virtualization at all. A business should list all the software applications they will be taking forward when migrating from a physical to virtual environment. They should contact the application's vendor or test the applications themselves, to see if they will work.



## Naming Conventions

Organizations usually have their own naming conventions in place to help identify servers, end-user desktops, etc. When a P2V migration takes place but physical machines still exist (which do not host virtual machines), it is often practical to name virtual machines in such a way, so that they are easily identifiable from their physical counterparts.

If this is done, then these changes will also have to be made at the network level, meaning an IT professional has much more work to do. If naming conventions are not changed, then the situation is effectively a 'plug and play' one, meaning that a virtual machine can be connected back on to the network without having to reconfigure any network setting (assuming of course that no connectivity changes have been made).

If a fast and easy migration is required, then naming conventions and network parameters should remain the same. If an organization has manageability in mind, then naming conventions and network parameters should be changed so that virtual machines are easily identifiable.

Once planning is over, the next step involves using a P2V virtualization tool. A P2V virtualization tool takes care of most of the migration process so that administrators don't have to. There are a number of different vendors and instances of P2V conversion tools but perhaps the most well known is VMware's Converter Tool.

VMware's Converter Tool is simple enough. It takes a snapshot of an entire machine's hard drive at the start of the conversion process. It then virtualizes this snapshot (converting it into virtualized files) and copies these over onto a new host machine.

Once the migration process is completed, there are a number of things that IT professionals should do before they actually give the virtual machine a 'green light' and let it go live.

## **4.2 Things to Do After Migration**

### **Disabling unnecessary Hardware within the Virtualized Environment**

When a physical machine is virtualized, it will still have hardware features such as USB ports, floppy disk drives, audio cards, etc enabled within its virtualized environment. If these hardware features are no longer required by your new virtual machine, then they should be turned off or disabled. This can both improve security and free up resources at the same time.

### **Starting-up the Virtual Machine without connecting it to the Network**

When starting up a virtual machine for the first time, its network card should be disabled or its interface should not be connected to anything. This is done to test the stability of the virtual machine when it is running as a stand-alone machine. A number of steps should be carried out while the network adapter is disconnected.

### **Testing critical or important applications**

Once a virtual machine has loaded up (while still disconnected from the rest of its network), important applications should briefly be run and tested to ensure that they work stably. If an application does not run stably then it should be stopped and further tested in a separate environment.

### **Rebooting the Virtual Machine multiple times**

This is done for two main reasons:

1. To ensure that all logs are clean after the migration process has been completed.
2. To ensure that the machine is able to warm boot (re-booting up after being switched on) and cold boot (booting-up after being switched off), without any problems.

### **Check and configure Server Parameters**

After all the above tests have been carried out successfully, the final step before going live involves checking network settings; for example, IP address configuration, subnet mask information and so on. If any changes need to be made, then they should be made at this stage.

Once all necessary changes have been made (if any), the virtual machine is ready to go live. You should first shut it down, then reconnect its interface and start it up again,

knowing that the migration process from physical to virtual has now been successfully completed.

It can be so easy for organizations to use a P2V tool to convert their physical machine into a virtual one without putting much thought into the process. However, without proper planning or testing, many unseen problems can suddenly pop up.

To ensure that a business's transition from a physical to virtual environment is as smooth as possible, they should carefully plan their transition from P2V and then fully test their virtual machine once it has been created. This will reduce the amount of problems that a business will face and make virtualization a very easy process.

### **4.2.1 A Simple How-To Guide**

This is a simple guide detailing the steps involved in the implementation stage of a migration process.

Careful and detailed planning should first be carried out before moving on to this stage. During the planning stage, administrators should be able to clearly identify what will be virtualized, for example what machines are going to be virtualized, whether applications are going to be virtualized, etc.

Once this has all been worked out, the implementation process starts. This stage involves actually ordering in any required hardware and then deploying the virtualization software onto these machines. There are three main steps during this process which are detailed below.

#### **Step 1: Order in the correct hardware if it's necessary**

If your existing servers and machines will be able to continue working efficiently in your new virtualized environment, then this is great. However, in many instances when a physical to virtual migration takes place, new hardware is required.

You should have already determined the specifications for the new machines that you require, during the planning stage. Now you need to purchase these from a suitable vendor. You should match servers from each of the different vendors and compare the prices, to ensure you get the best deal.

One of the best hardware solutions, when migrating from a physical to virtual infrastructure, is converting your existing hardware infrastructure to a blade server (if it isn't already).

Blade servers are servers which can easily be added or removed from an enclosure and they are independent from storage and network connections. The obvious benefit of having a virtualized blade server environment is that it offers much more flexibility than a standard virtualized server environment.

## **Step 2: Select your virtualization software**

Again, this is something that should have been considered during the planning stage but some aspects are automatically decided. For example, you wouldn't select a virtualization application that only works on Linux-based operating systems, if all your servers were Windows-based. Selecting a virtualization solution may seem tricky but there are a number of factors that will influence your decision. These include:

*Price* – How much does the software cost or is it open source?

*Type* – Is it a bare metal virtualization solution or a hosted one?

*Compatibility* – If it is hosted will it be able to run within our host operating environment? Also can it run Windows/Linux/Mac OS X operating systems? Finally will our existing applications run on it?

*Hardware requirements* – Will the virtualization software run on x86-based operating systems? What are the minimum hardware requirements that it supports?

Using the above factors, along with the research you carried out during the initial planning stage, you should be able to come to a conclusion and select the virtualization software that is right for you.

You should also note that some operating systems will come bundled with their own virtualization software. For example Windows offer their Hyper-V virtualization solution on Windows Server 2008 and Ubuntu Server Edition comes bundled with its KVM (Kernel-based Virtual Machine) virtualization solution.

When buying new servers, it can sometimes be very cost effective to use the virtualization software of the operating systems that will be installed onto them, rather than buying separate stand-alone virtualization software. This saves money because organizations do not have to buy any additional virtualization software for their new or existing servers.

However, operating system virtualization solutions are hosted solutions (because they are not directly installed onto hardware), therefore they will not be as efficient as bare metal

virtualization solutions. If a business wants maximum efficiency from their machines, then they should install bare metal virtualization software.

VMware ESX Server is one of the most popular bare metal virtualization solutions, although it is quite expensive when compared with other virtualization solutions. It works when installed directly on to server hardware in place of an operating system. Because it interacts directly with hardware just as an operating system would, it is far more efficient than hosted virtualization solutions are.

Finally, select a physical to virtual (P2V) conversion tool. Although migration can be carried out without a conversion tool, it is not recommended because it is a tedious, time-consuming task and people are prone to making errors. A P2V conversion tool streamlines the conversion process and takes care of the small tedious things so that you don't have to. Even better is the fact that most major vendors of virtualization software offer free conversion tools, meaning that you have no excuse not to use one.

### **Step 3: Carry out the actual virtualization process**

Once all of your hardware has been ordered in and you have chosen your virtualization software, you can begin the conversion process using your P2V conversion tool. The conversion process is briefly outlined below.

1. The conversion tool takes a snapshot of the existing system.
2. The snapshot data is encapsulated in to a virtualized form. For example, virtual machine files are created from the snapshot information and all of the data from the snapshot is encapsulated within a virtual hard drive.
3. These virtualized files are then transferred over to the new hardware, or the existing machine's data is overwritten with this new virtualized information (depending on whether new or existing hardware is being used as part of the new virtualized infrastructure).
4. The virtual machine is booted-up while it's disconnected from the network and thorough testing is carried out.
5. Once testing has been successfully completed, any settings or parameters that need to be changed or configured are done at this point.
6. Finally the virtual machine can go live and is connected to its network. The old physical machine can be disconnected (if it is no longer a part of the new infrastructure) and the migration process is now completed for that machine.

These steps should be carried out for all machines that are going to be virtualized.

## 4.2.2 Further Migration Considerations

As the market for virtualization technology has grown, many different vendors and technologies have emerged. Each different virtualization technology is likely to slightly differ in the way that it is setup or implemented. However, there are things that you should do no matter which virtualization technology you use. The top ten things that you should do or consider when migrating from a physical to a virtual environment are detailed below.

### 1. Considering management costs and which management tools to use

The number one reason why businesses migrate from a physical to a virtualized infrastructure is because of the large amount of money that it will save them in terms of hardware, power, space, etc. This is true, because virtualization will save money in these areas, but a lot of the time businesses fail to work out what the management costs for a virtualized environment would be.

If the management costs for a virtualized solution will balance out the savings that a business initially makes, then it is not worth it for a business to carry out virtualization. Management tasks include updating an operating system and its applications, applying patches, backing up data, installing new applications and so on.

However, luckily for most businesses, there are virtualization management tools which can significantly reduce the costs of administering a virtualized environment. Most well known vendors such as Microsoft and VMware will provide management tools with their virtualization software (sometimes even for free). These will often have features such as P2V (physical to virtual), V2V (virtual to virtual) and V2P (virtual to physical) conversion. Conversion tools make the migration process or the management of entire virtual machines much easier.

When migrating from a physical to a virtual environment, businesses should not forget to include management costs during their planning stage so that they can see whether migrating actually would save them money. Finally, a business should consider purchasing virtualization management tools which make administering a virtualized environment easy, saving them money in the long term.

### 2. Deciding on the type of virtualization solution to use

Although this step may seem simple enough, it is actually quite important and should not be overlooked. There are three main categories of virtualization solutions that a business can choose. These include bare metal virtualization, hosted virtualization and OS virtualization.

Bare metal virtualization involves installing virtualization software directly on to a machine or server's hardware, just like an operating system would be installed. Virtual machines can then be deployed under this virtualization software, without the need for a host operating system in-between. The main advantage of using a bare metal virtualization solution is that it is currently the most efficient virtualization technology out there, because it interacts directly with hardware. If businesses require high efficiency or near-native performance from their virtual machines, then this is the virtualization software they should use.

The next type of virtualization is hosted virtualization. Hosted virtualization involves installing virtualization software, in the form of an application from within a host operating system environment. Virtual machines can then be run via this software. Basically this means that for a virtual machine to interact with its hardware, it must first go through its virtualization application and then through its host operating system.

In a bare metal environment, all a virtual machine has to do in order for it to communicate with its hardware is interact with its virtualization software. This means that a bare metal solution is much more efficient than a hosted solution. However, hosted virtualization solutions are much cheaper than bare metal virtualization solutions. Therefore, if a business is only going to be carrying out virtualization for a few of their physical machines, then they should consider using a hosted virtualization solution to accomplish this.

The above two solutions have one major pitfall. Instead of decreasing the amount of administration that is required, they increase it. The above two solutions make it so easy to create virtual machines, that businesses often create more than they need. This means that more management is required because there are now more software environments than previously and multiple instances of the same OS or even different ones (making it that much harder to administer a virtualized environment).

The third virtualization solution (OS virtualization), effectively resolves the above problem. OS virtualization involves running multiple instances of the same operating system but under a single operating system (and a single kernel). The disadvantage is that different operating systems cannot be used on the same physical machine. However, because everything is run under a single kernel and operating environment, all an administrator has to do is manage this top level environment, which greatly decreases their workload.

As you can see, there is more than just one virtualization solution that businesses can use when migrating from a physical to virtual environment. Each solution has its own benefits, depending on what a business's requirements are.

If a business requires efficient performance, then they should use a bare metal virtualization solution. If they are looking for a cost effective solution, then using a hosted

virtualization application would be beneficial to them. Finally if a business wanted a solution that was easily manageable and could save money on management costs, then an OS virtualization solution would be the solution for them to implement.

### **3. Considering performance requirements**

Performance is always an important requirement no matter what type of system or infrastructure is being implemented. In terms of virtualization, the factor that will affect performance the most is processing overheads.

The processing overheads for virtualization depend on what type of solution is being used. Virtualization solutions from different vendors may also perform at different efficiency levels. The performance of different virtualization solutions can be as efficient as having only 1 per cent of processing overheads (providing near-native performance), or they could be as slow as having 60 per cent of processing overheads.

If critical or real-time applications are going to be run in a virtualized environment, then a virtualization solution should be chosen that can perform at near-native speed. However if applications are not resource intensive, then they can be run on less efficient virtualization software.

When requiring a highly efficient virtualization solution, bare metal virtualization solutions or OS virtualization solutions are the most efficient. Hardware-assisted virtualization technologies or alternatively Para-virtualization technologies also improve efficiency. Hosted virtualization solutions are not as efficient because there is an extra layer of communication which increases processing overheads. However, if virtualization is not required for high performance or real-time applications, then a hosted solution is a reasonable choice because it is reasonably cheap and relatively easy to implement.

### **4. Virtual server density**

This simply refers to how many virtual machines can run or 'fit' on a physical server without a significant decrease in performance for all virtual servers. Having physical servers which can run multiple virtual machines is clearly beneficial to a business, as it allows them to save money by not requiring as much physical hardware as they did before.

However, some virtualization solutions may have limitations and businesses should take this into account when selecting a virtualization solution. There is nothing worse for a business than ordering in new high performance physical servers to run five virtual servers for example, only to realize that the virtualization software they purchased can only run two virtual servers.



Hardware performance is another factor that will affect a server's density level. There are only so many virtual machines that a processor can deal with before it comes under serious strain. With advances in hardware-assisted virtualization technology, physical machines can now accommodate more virtual machines today, than ever before.

If a business needs to run multiple virtual machines on a single physical machine, then OS virtualization is probably the best solution for them. This is because an OS virtualization solution controls all of its separate virtual environments under a single kernel. This creates a much smaller and manageable workload for a processor (when compared to other virtualization solutions which have multiple kernels), meaning that an OS virtualization solution can run many more virtual machines than other virtualization solutions can.

## **5. Virtualization software platform support**

Although this may seem obvious, not all hardware platforms are supported by all virtualization software. In fact, the hardware formats that a business use can often decide the virtualization software that they will ultimately use.

Hardware-assisted virtualization solutions and Para-virtualization solutions are the least compatible because they must support the hardware that they reside on. OS virtualization and hosted virtualization solutions are the most compatible, because these solutions run under their host operating system, which already supports its physical machine's hardware.

Different virtualization solutions will also differ in the type of hardware that they can emulate. Some virtualization solutions may only be able to emulate 32-bit processors, while others may only be able to emulate a single processor, or others multiple processors.

Based on what businesses require or what hardware format they use, there is a different virtualization solution out there for each one of them and only they can decide which virtualization solution is compatible with their physical infrastructure or goals.

## **6. Virtual server migration and management strategies**

Managing a virtualized environment is different from managing a physical environment. The greatest difference between managing a virtualized infrastructure when compared to a physical infrastructure is that a virtualized infrastructure is hardware independent. This means that a virtual server is not stationary, rather it can move from one physical server to another physical server.

This provides many benefits. For example a business can:

- Easily move their virtual machines onto more powerful hardware;
- Avoid downtime when a physical machine has to be upgraded or repaired by moving its virtual machines onto another physical machine; and
- Effectively make their virtual machines portable. For, example MojoPac can create portable virtual machines on any kind of USB drive.

A business should select a virtual machine management solution based on their needs and budget. If a business requires maximum uptime with an absolute minimum of downtime, then they should select a high end virtual machine management solution. A high end management solution is likely to be costly (due to software costs and redundant storage costs), however it will make it easier for an organization to migrate, clone or even back up their virtual servers, in case there are any problems.

Cheaper management solutions are available, however they often do not guarantee uptime at all times. Again, management is an important issue to think about when an organization is going migrate from a physical to virtual infrastructure. Although it won't initially affect a business, eventually their management solution will come into play.

## **7. Deciding on the level of isolation use**

Different virtualization tools provide different levels of isolation for virtual machines. Bare metal virtualization solutions provide complete isolation for their virtual machines. This means that no virtual machine can see or interact in any way, with any other virtual machine that resides on a single physical machine (unless of course they are connected to each other via a business's network).

Hosted solutions are effectively the same, although some hosted solutions do not completely isolate virtual machines from one another. For example some hosted solutions allow separate virtual machines to communicate with each other, via their hosting environment (host operating system environment).

If a business wanted to completely isolate their virtual machines from one another, then they would use either a bare metal virtualization solution or a hosted virtualization solution as detailed above. However, there is another solution for businesses: this is known as application virtualization.

Application virtualization involves running applications in isolation from one another in their own virtualized environment. They are all run under the same operating system or operating environment. If a business's sole purpose for converting their infrastructure from a physical one to a virtual one was to separate applications from one another, then an application virtualization solution would be a much cheaper solution for them than a full virtualization solution.

## 8. Customizing a virtualization solution based on its intended purpose

Virtualization solutions are implemented for many different reasons but there are three that are most common. These reasons are, for testing and development purposes, to consolidate server resources and for disaster recovery purposes. Based on what a virtualization solutions purpose is, it should be customized accordingly.

### *Testing and development*

Virtualization is commonly used to create isolated testing and development environments which won't interfere with or damage a businesses existing infrastructure. If an application is being developed for many different operating systems, then a bare-metal or hosted virtualization solution is the best solution to use.

This is because using these solutions, multiple virtual machines with different operating systems can be created. The added advantage of using a hosted virtualization solution is that it is cheaper and allows IT professionals to switch between their testing/development environment and their actual physical or normal workspace.

If an application is being tested for a single operating system, then a good solution is using an application virtualization program. Again this is cheaper than a full virtualization solution and it allows an IT professional to test an application in the normal working environment, but in complete isolation. This means that their host operating environment is completely safe and can continue running normally, no matter what other applications are doing.

### *Consolidating Service Resources*

Here, virtualization is used so that multiple virtual servers can be aggregated on to a single physical machine, better utilizing its resources. The virtualization solution that a business requires in this case is one that is efficient and produces low processing overheads.

In this case a bare metal virtualization solution is the best choice. This virtualization solution provides the greatest performance-to-processor overhead ratio when compared to all other virtualization solutions. Additionally, hardware assistance is preferred because this provides even greater efficiency when compared to just a fully software-based virtualization solution.

Finally, if all physical servers use the same operating system, then an OS virtualization solution is something that should also seriously be considered. OS virtualization technology is highly efficient, because everything is controlled by a single kernel, meaning far less emulation has to take place.

### *Disaster Recovery*

In some cases, businesses just use a virtualization solution as part of their back-up and disaster recovery solutions. Backing-up data or providing redundant devices in case of failure, is often a time consuming and costly process.

With the advent of virtualization technology this has changed and businesses have picked up on this, using virtualization to make their backup solutions a lot more effective. Virtual machines can easily be created, copied and even moved from one physical machine to another.

This means that administrators never need to worry about losing data caused by a fault or due to hardware failure for example (as long as they take regular snapshots of course). If a physical machine does fail, then all an administrator has to do is copy over the most recent snapshots of their virtual machines, onto another physical machine and things can return to normal.

The best solution to use for backup purposes is a hosted virtualization solution, along with a good virtualization management tool. Using a hosted virtualization solution, virtual machines files can easily be copied or backed-up from within their host operating environment.

## **9. Check licensing requirements**

In terms of licensing, most software vendors consider a virtual machine to be the same as a physical machine. This means that you will still need licensing for all the different instances of an operating system or applications that you install, whether they reside on a single physical machine or not.

Certain vendors' terms may also not allow their operating system to be used on a virtual machine (although more and more vendors are supporting virtualization by the day). The licensing terms for Windows Vista home editions (both basic and premium) initially disallowed these operating systems from being installed on a virtual machine (although other versions of Vista were allowed to be installed on virtual machines). However this has now changed and all versions of Vista are now allowed to be used in a virtualized environment.

Different versions or editions of an operating system may only provide a certain number of licenses for virtualized installations. Businesses should always check that the operating

system they are installing onto their virtual machines supports virtualization and provides them with the number of licenses they need.

## **10. Use templates for automated system deployment**

One of the great things about a virtualized environment is that it is much easier to provide and deploy new machines than it is with a physical environment. Below are simple steps to create a virtual template.

*Step 1:* Create a virtual machine and install the desired operating system on to it.

*Step 2:* Once the operating system has been installed, you should install all patches and any applications that will be common to all virtual machines. Any configurations that are common to all virtual machines should also be made at this stage. Applications that are limited to only a few virtual machines will be installed later on.

*Step 3:* Generalize this virtual machine so that it appears as a template and will not interfere with existing machine names or IP addresses. For example it shouldn't have a name or IP address that will conflict with any other machine's name or IP address.

*Step 4:* The template should now be stored or archived in a secure location where it will not be changed or modified in any way. The same precautions used to secure normal data, such as password protection or encryption should be considered to help secure virtual machine templates.

*Step 5:* At this stage, virtual machine templates can be deployed to physical machines ready for use. This can either be done manually or automatically using software. Doing this step automatically is the preferred method.

*Step 6:* Any new updates should be installed onto the virtual machines now in operation. Any applications that are only going to be used by a few machines should also now be installed on to these machines.

Using templates can save IT professionals a lot of time so they can use their expertise elsewhere on more important things. Any business that uses some form of virtualization should seriously consider using virtual templates.

### **4.2.3 Risks Associated with Virtualization**

Although virtualization provides businesses with many benefits, there are some pitfalls. Detailed below are ten of the most common pitfalls and how to avoid them.

## **1. Not being able to see or manage virtual machines**

Physical machines are easy to manage and keep track of because they physically exist. However virtualized or software-based machines can actually be hard to keep track of. For example it is easier for an organization to count how many physical servers they have but it is much harder for them to count how many virtualized servers they have (without the proper tools of course).

Businesses should use some kind of virtualization management software to help them keep track of their virtualized environment. Using a virtualization management tool means a business can easily see how many virtual machines exist on their physical machines, which virtual machines are active and which are inactive.

## **2. Having difficulty keeping track of what each virtual machines role is**

Again because virtual machines are not physical and are software-based, it can sometimes become difficult, keeping track of what each virtual machine is doing or what each virtual machines role is.

Businesses can sometimes even lose track of which physical machines are hosting critical virtual machines, such as file servers or internet gateway servers, etc. If this happens, a business could mistakenly switch off a physical server which is hosting critical virtual machines that require maximum uptime. A simple mistake like this can cause a lot of problems for a business.

The solution to the above problem is not only using good management tools again but also grouping similar virtual machines together, so that confusion can be avoided. For example all critical virtual machines could be placed on a single server, while less critical virtual machines could be placed on different servers, depending on their roles.

## **3. Not having enough security on a physical host machine**

Although virtual machines provide certain security benefits to a business, the physical machines they are hosted on still have the same vulnerabilities as before. In fact, in a certain way, more damage can be caused to a virtualized infrastructure via the compromise of just one physical machine, than can be caused to a normal physical infrastructure.

In a physical environment, if one machine is corrupted by a virus for example, then only that machine is affected. This is assuming of course that the virus is not transferred to other machines via a network or through some other means.

In a virtualized environment, if a host machine is infected by a virus that corrupts data, then all of its virtual machines will also become corrupted because these virtual machines are also data. The damage caused can be huge, but fortunately virtual machines are easier to back-up and as long as a business takes regular snapshots, they should be relatively safe if their virtual machines are ever corrupted.

Business should not take for granted the extra security that using a virtualized infrastructure brings. They should ensure that both their physical and virtual machines are secured with the latest software patches, firewalls and anti-virus software.

#### **4. Not knowing who manages what**

A physical infrastructure is relatively easy to manage and everybody knows what their role is. For example, hardware professionals manage the physical machines, while software professionals manage the software on these machines.

In a virtualized environment, employees can lose track of their roles. For example, virtual machines are software, so are they now managed by the software professionals or do the hardware specialists continue to manage them?

This kind of confusion can lead to jobs being overlooked or important tasks not getting completed because each person thinks that it is the other person's job to complete certain tasks.

To avoid this type of confusion, when the new virtualized infrastructure is implemented, each employee should be clearly briefed on what their role is in managing the new infrastructure, how it is the same as before and how it may differ from their role before.

#### **5. Not controlling or managing the deployment of virtual machines**

Virtual machines have the ability to move from location to location or from one physical machine to another with ease. Although there are many benefits such as the ability to keep virtual machines running when physical machines fail, there needs to be some way to keep track of their movement.

Things can quickly get confusing in a business environment when virtual machines are constantly being moved from one physical machine to another. A process should be implemented where it is made clear to all employees which physical machines their virtual machines are located on.

#### **6. Not keeping track of resource utilization or applications**

Once a virtualization solution has been successfully implemented, businesses often don't keep track of their resource utilization. For example, during the initial planning stage, a business may have worked out that they can run five virtual servers per physical machine efficiently. However, when the virtual machines are actually running, they will each have their own applications that will increase a processor's workload.

Some applications may bring a physical machine to a complete halt, affecting the performance of all other virtual machines. If businesses do not keep track of which virtual machines are slowing down a system or what applications they need to deploy somewhere else (for example on a machine with less virtual machines), then they are not making the most effective use of their physical resources.

## **7. Not preventing virtual machines from being copied by unauthorized personnel**

Virtual machines are so easy to copy, that practically anyone can copy them. This poses a problem when multiple individuals use the same physical machine. Although their data will be isolated from each other, there is nothing stopping them from copying another user's virtual machine onto an external drive, so that it can be hacked into at a more convenient place, for example.

There are a number of precautions that businesses should take to prevent their virtual machines from being copied by unauthorized personnel. Firstly they should disable all ports such as USB ports if they are not required. A business could even use physical machines that have no USB ports to host their virtual machines.

Next they should restrict what a user can see at the host operating system level. To accomplish this, separate partitions could be created or data could even be encrypted so that restricted data is unreadable.

However, the best way to prevent virtual machines from being copied in a shared physical environment, is implementing a desktop virtualization solution that uses thin-client technology.

What this means is that every user's virtual machine data is stored centrally on a server (or on multiple servers). The only way a user can access their virtual environment is by inputting their correct login information, using thin-client technology. This completely stops virtual machines from being copied by any unauthorized personnel because only the individuals, who provision the servers on which the virtual machines are stored on, can actually see them.



## 8. Forgetting licensing costs

During the initial planning stage, it is easy for businesses to forget about licensing costs. For example, if a business was using five physical machines to host twenty five virtual machines, it can be easy to make the mistake of thinking that only five licenses are required.

However, most commercial operating system vendors see virtual machines as being exactly the same as physical machines. In the scenario above, this means that a business would actually require thirty licenses. This could be overlooked when dealing with the other bigger aspects that a virtualization solution provides businesses with.

One great way to save money on licensing costs is to use an open source operating system. If a business is going to have many virtual machines running within their environment, then the most feasible solution for them would be to use an open source operating system such as Linux.

## 9. Hidden costs

Similar to the licensing issue as detailed above, businesses often overlook things that may actually end up costing them money, rather than saving money. Licensing is one of these things. Another is hardware costs. To a business it may look like they are saving money on hardware by running five virtual servers on a single physical machine, however this may not be the case if they have to order in new hardware.

Management costs are also something that businesses can overlook. If a business has a single physical server, which contains five virtual machines, it's important to include the management costs for these five machines. They see a single physical machine and believe only one machine needs managing. However this is not true, five machines still need managing (or six if you include the host machine).

A business should always carefully scrutinize their costs and savings when migrating to a virtualized environment. They should consider all costs including those that may be less obvious. If the result is that they are spending more money on the virtualization process than they are actually saving, then it is probably not worth migrating for that business at that particular time.

## 10. Overloading or underutilizing physical machines

Sometimes organizations overestimate their physical server's ability to handle multiple virtual machines. They keep creating more and more virtual machines on a single physical machine to save money. However, eventually the physical server will reach its limit. When

this happens, application productivity will go down, and this can actually cost a business in the long run.

Other times, a business may underutilize their physical machines instead of over-utilizing them. For example they may only run three virtual machines on a physical server that is capable of hosting five virtual machines efficiently. Again, this is costly for a business because they have to purchase or use additional hardware that is unnecessary.

Like most things, finding the right balance is the key to having an efficient and money-saving virtualization solution. A virtualization solution should utilize around 80 per cent of a machine's physical resources. Any less and resources are being wasted, any more and the physical machine will probably be under strain, decreasing its overall performance.

#### **4.2.4 Problems Associated with Virtualization**

Described below are some of the problems that may be faced by businesses and IT professionals when a virtualized infrastructure is used.

##### *Uptime requirements increase even more than before*

Although storing multiple virtual machines on a single physical machine provides businesses with many advantages, this practice leads to even more stringent uptime requirements than were needed before.

A physical machine must meet the requirements for its most critical virtual machine first. If one virtual machine requires 99 per cent uptime, and a physical machine's other virtual machines only require 90 per cent uptime, then that physical machine must still remain functional 99 per cent of the time.

This also means that if it is harder to schedule downtime for upgrades or maintenance because you have to find a period of time where none of the virtual machines are being utilized. The more virtual machines there are on a single physical machine, the harder it becomes to schedule downtime and the higher the uptime requirements for that physical machine become.

##### *Bandwidth related issues*

In a physical environment, all physical machines that are connected to a network use their own dedicated NIC (network interface card) to connect to the network. However, in a virtualized environment, multiple machines might share the same NIC. This creates a

scenario where a single physical machine is connected to a single port on a switch, router or other networking device (because there is only one NIC).

In a normal physical environment this works fine, a single port and interface should be able to accommodate the bandwidth requirements of a single physical machine. However, on a host machine supporting multiple virtual machines which are all using a single NIC and port, bandwidth related problems are often created.

Due to the large volume of data that is created by multiple virtual machines, a network can often be brought to a standstill. The individual ports on networking devices such as routers or switches were never designed to accommodate the bandwidth requirements of more than one machine.

One way to reduce the problem detailed above is by having a physical machine with multiple network interface cards. At least this way, each virtual machine can be assigned their own network interface and switch or router port. However, once again there is only so much space inside a physical machine and at some point even the process of installing multiple network interface cards within a single machine will no longer be feasible.

### *No support for critical applications*

Although thorough testing should be carried out during the planning stage, for whatever reasons, business may sometimes forget to test certain applications or overlook things. Once their new virtualized infrastructure has been implemented, they then find out that some of their most important applications no longer work in their new virtualized environment.

The consequences of this are huge. Once a business realizes that some of their most important applications no longer work, it is difficult to go back to their old infrastructure and also very costly, because they would have to effectively pay twice as much to do go back.

Although the concept of virtualization has been around for many years, modern day virtualization technologies are still relatively new. Some businesses are avoiding compatibility issues by just not migrating at all (at least not yet anyway).

By waiting a few years, a lot of the 'teething' problems that virtualization technologies currently have, should be smoothed out or even eliminated. As virtualization technologies improve, not only should they become more efficient, but they should also become more affordable.

If a business is going to implement virtualization technologies now, they should ensure that their current software is able to run on their new virtual machines.

## ***5 Managing a Virtualized Environment***

### **5.1 Support Issues**

When talking about support issues in a virtualized environment, there are two possible meanings. Firstly we have technical support, which is basically the personnel who support virtualization technology and deal with this technology in a business environment. Secondly we have software related support issues. This refers to the compatibility between different virtualization vendors, operating systems and virtualization software and finally software applications and virtualization software.

Detailed below are the two types of support related issues.

#### **Technical Support**

When a business implements virtualization, there are many things that can easily be overlooked or forgotten about. Technical support is definitely one of these things. When a business runs software or applications in a normal physical environment, technical support will often automatically be provided as a part of their purchase.

With the advent of virtualization technology, things are not as straightforward. For example, if a certain vendor's application is run in a virtualized environment, even though it was not originally intended to, will that vendor still offer technical support? This is an important question for businesses that they must find an answer to.

Traditionally, software application used to run within a single environment. This meant that technical support staff could all follow a straightforward script when troubleshooting.

This was until the client/server model became well established. The client/server model allowed software to be distributed and run from many different locations. This now made things a little bit harder for technical support staff, because there were now multiple different scenarios for them when troubleshooting and a single problem could have been caused by a number of different things.

Virtualization further adds to this complexity. The more programs or software applications that are running on a single physical machine, the more potential ways there are for a

problem to occur. This obviously makes the job for technical support staff much more difficult.

Also, because most software vendors designed their software to be used in a physical environment, this is the kind of environment that the application would be tested in. It is unlikely that they would have carried out tests in a virtualized environment, therefore most software vendors may be unsure whether their software could run stably in a virtualized environment.

This also means that a software vendor's technical support staff would be unlikely to have the knowledge to offer support to a business that uses virtualization technology. One of the worst things that can happen to a business is that one of their critical applications stops working after having migrated to a virtualized environment. They then go to contact the technical support staff for that particular application, only to find out that technical support does not offer support or just doesn't have the knowledge to offer support for that businesses virtualized environment.

The main reason why the above scenario or similar scenarios occur, is because technology suppliers follow what is known as an 80/20 rule. This means that they only test their technology on the most commonly used platform. This usually consists of around twenty percent of all possible platforms. They do this, in the hope that this will fulfil the needs of around eighty percent of all customers.

New technologies are obviously not commonly used because they are still emerging. This means that software vendors often don't test their software on new platforms until they become widely used. As virtualization becomes more widely used, more vendors will start designing and testing their applications to be run in virtualized environments.

Until then however, businesses should always confirm with their software vendors whether they offer technical support for their virtualized environment. Finding out this information beforehand, rather than at a later date can save businesses a lot of hassle and reduce downtime caused by these problems.

## **Software Support**

Software support refers to which software is supported by which platforms; what applications can be run in a particular environment and so on. Firstly, we have support at the hardware level. This level will determine the fundamental virtualization software that a business uses, whether it is a bare metal solution or a hosted solution. If a business has physical machines which use a 64-bit processor, then they will need to deploy virtualization software that offers support for 64-bit processing.

The next level of support is at the host operating system level, assuming of course that a business selected a hosted virtualization solution, rather than a bare metal virtualization solution. If a business operates using a Linux-based infrastructure, then the virtualization software that they deploy should support Linux-based distributions.

The next level is at the guest operating system level. Most businesses will probably want to use the same operating systems that they were previously using on their physical machines. However, in some cases a business may also want to run other virtualized operating system environments. Again, a business should confirm that the virtualization software they are using supports all the guest operating systems that they require.

Finally we have the application level. This refers to which applications can run within a particular virtualized environment. Some of these applications are the same as before. For example you can't run Windows executable files on a Linux-based virtual machine (unless it is using application virtualization software).

However, some applications which did previously work within their native operating environment on a physical machine may not actually work within a virtualized environment that uses the same operating system. Businesses should always confirm with their software vendors whether certain applications will work in a virtualized environment.

Microsoft recently created (since June 2008) what is known as "Microsoft's Virtualization Support Policy". This allows virtualization software vendors to test and validate whether Windows Server 2008 will run on their virtualization software without any problems.

As more software vendors offer support for their applications to be used in virtualized environments, most software that is currently supported by physical machines should eventually be supported by virtual machines.

Virtualization technology has come a long way and it is only a matter of time before virtualization solutions become as common as physical solutions and offer the same level of software support.

## ***5.2 Measuring Capacity and Performance***

Whenever a new technology emerges, it creates a whole host of new challenges including the challenge of monitoring that new technology's performance. This has been true for all of the other major advancements which are commonly used by businesses today. Virtualization is the next major advancement in relation to computer technology.

New technologies such as Intel and AMD's hardware-assisted virtualization solutions along with many others advancements in the area of virtualization have caused more and more businesses to migrate from a physical to a virtualized infrastructure. This has completely changed the way a business operates and how they now manage their new and improved infrastructure.

### **The problems that businesses face when they cannot efficiently manage their virtualized environment**

Businesses that cannot efficiently monitor or manage their physical (or virtual) machine's performance are bound to face problems sooner or later. Without being able to see detailed information about what is going on inside their physical machines, businesses have no idea what the reasons behind their problems may be.

For example, businesses may have five virtual machines on a physical server that is capable of hosting ten virtual machines. A business would expect the physical server to run efficiently and its performance to be at a high level, because only fifty percent of its total resources are being utilized.

However, if suddenly that physical machine no longer performed as efficiently as it should be performing, then without the use of any performance management tools, a business would find it very hard to tell what the cause of that problem was.

Of course there could be many reasons why a physical machine is no longer working as efficiently as it should be. In the scenario described above the problem is not the physical machine itself (because it is only under 50 per cent load); it is actually other external factors which bottleneck performance, such as network link speeds, storage access speeds and so on.

Without the use of good performance management tools, it would be a lot harder and take a lot more time for a business to work out that it was actually external factors limiting the performance of their machines, rather than the actual machine's hardware being the limiting factor.

### **Monitoring and managing performance**

When a business manages their virtualized environment, the best possible solution for them is to manage their virtual machines across an entire domain, rather than individually. This is beneficial because it saves businesses both time and money on management-related issues.

The type of management described above is known as “Total system performance management”. This type of management effectively groups together all of the individual elements of a virtualized infrastructure (such as virtual desktops, virtual servers, etc), into a single container or ‘black box’.

The performance and capacity related issues (along with other issues), are then measured for this entire ‘black box’. Performing management this way can ensure businesses quickly get an idea of how efficient their virtualized environment is as a whole, or how effectively their storage area network (SAN), is being accessed by all of their virtual machines.

As you can see, a business’s management strategy can be made a lot more effective and easier by aggregating the individual elements of that business’s infrastructure into a single container. This container or ‘black box’ can then be monitored and managed by IT professionals.

One final strategy involves creating their own metric or to identify a metric which can measure the performance of their virtualized environment as a whole. One example of a complete system metric is I/Os per second (IOPS). This can measure the total usage of a system showing how efficiently it is reading and writing data.

When managing a virtualized infrastructure, it is clear that the total system approach is the best way. It can group together the individual layers of a virtualized infrastructure such as virtual machines, network connections and storage devices measuring their overall performance and efficiency.

The use of metrics further simplifies this as data can be easily presented to IT professionals in a simple and understandable way. The best way for a business to implement a management solution like this is by firstly identifying what metrics a business will use to measure their performance. Then software management tools that can measure performance across an entire domain should be implemented. These can provide IT professionals with information that relates to the performance or capacity metrics that they are using.

### ***5.3 Contracts and Agreements such as Licensing***

Licensing is one of the things that can easily be overlooked by a business when they are migrating from a physical to a virtualized environment. However, it is quite an important aspect for a business to consider, because it can determine what software a business is allowed to use with their new virtualized infrastructure and how they will be priced for using this software.



Many traditional licensing models are now becoming completely useless in a virtualized infrastructure, therefore new licensing models have to be created or old ones adapted to meet the needs of today's technology.

Most software for servers are now licensed per CPU and pricing models for software are being created, which are based on how much memory or how many virtual cores exist on a single physical machine.

Licensing is slowly being changed from the traditional way, in which licenses were given for the number of copies of software that a business had, to the new way, in which licenses are handed out based on the number of processing cores or amount of hardware that a business has. This is a lot easier to keep track of hardware.

Software can easily be changed or copied. This means that it can be difficult for a business to meet their licensing requirements when licensing is based on software. This is especially true in virtualized environments, where it is that much more difficult to keep track of software.

Hardware, such as the total number of processors that a business has or the total amount of memory that their machines have, are unlikely to change, at least not frequently. This means that it is much easier for a business to meet their licensing requirements and they don't have to keep track of every instance of a software application that they have installed onto their machines.

It is much easier for software vendors (and better for businesses) to create a licence based on the amount of cores that a machine will have (whether they are virtual or physical), rather than per copy of their application for virtualized infrastructures. For example a software vendor could create separate licences for machines that had a single processor, two processors, three processors and so on. A business could then purchase the license that they required based on how many processor cores each on of their machines had.

If a business owned servers which had two processing cores, then they would purchase a licence for two processing cores. This business can now use as many instances of the software that they purchased, that their physical machine will allow. This means that they can run multiple virtual machines on that physical server, without having to worry about how many virtual environments the software application is running in.

This is beneficial to businesses, because they can save money on licensing when using a virtualized infrastructure. However, it is also beneficial to software vendors because they know that their licensing terms will not be broken (unless the number of physical

processors that a machine has increases, but this is not something that happens frequently).

Licensing issues can also affect the choice of virtualization software that a business purchases. For example, some virtualization software vendors may allow businesses to run more virtual machines than others, on the same amount of hardware. Businesses will be inclined to go for the software virtualization vendor that gives them the best deal for their money.

When thinking about migrating from a physical to a virtualized environment, businesses should not forget about their existing contracts and agreements. Making changes to their existing contracts and agreements (which is a very likely scenario when virtualization is carried out), can actually incur additional charges for a business.

Sometimes this cost can be so high that it can neutralize the cost-saving benefits of virtualization. It would be better off for some businesses to wait for their existing contracts and agreements to expire before migrating from a physical to a virtualized environment.

A business should have a good reason to justify the migration from a physical to a virtualized environment. Businesses should not carry out virtualization for the reason of saving money alone. They should carry out a virtualization process for the actual benefits that virtualization will provide, such as improved flexibility or for increasing their machine's storage capacity via storage virtualization. This will ensure that they are saving money in the long run. It's important to have a justifiable reason for carrying out virtualization, as well as expecting some cost-saving benefits, because unimagined problems (such as contract and licensing disputes) can end up costing a business more than initially forecast.

## **5.4 Organizational Considerations**

When carrying out virtualization, many changes are made to a business's infrastructure along with the way they now operate. Changes are also made to the way employees work and the roles they have in the new virtualized environment.

### **A change in both the ways and times that people work**

The first major change is that departments, teams or individuals who wouldn't normally share their machines or resources must learn to work together with the other areas of their business. Virtualization means that many different environments can exist on a single machine; so users who used to have their own single workstation are now likely to share this machine with other workers.

This has a number of consequences for businesses. Firstly they must ensure that workers can only access or see what they are supposed to. Also, each worker needs to be allocated a timeslot for when the physical machine on which their virtual machine exists, is free. This means that some employees may be working at completely different times than they used to.

One way to keep things as similar as they were before is by using desktop-based virtualization. Using this solution, virtual machines are stored on a server. They can then be accessed using any thin client and at any time. This means that most employees can continue working as they did before and at the same times as they did before.

### **Who is in control, what skills are required, etc?**

When virtualization is implemented, certain questions need to be answered by businesses. Some of these questions include: Who is in control of the new infrastructure? What new skill will be required by employees, if any? Are old departments still required or can they be downsized?

The first question stems from the fact that a virtualized infrastructure is difficult to administer centrally. Different resources are scattered across various servers and because of this, more than one administrator is usually required in a business that has a virtualized infrastructure (of course the number of administrators required also depends on the size of a business).

One potential strategy that businesses can use, is assigning administrators to control the different aspects of a virtualized infrastructure. For example, some administrators may only be assigned to administer physical machines, while others have the task of monitoring and managing virtual machines. This means that businesses have to re-train some of their personnel to work with the business's new virtualized infrastructure and virtualization software.

This leads on to the second question, what new skills will be required by employees? In most cases, administrators and IT professionals have a wide enough knowledgebase to work in a new virtualized environment without facing any problems. However, businesses should still offer training and train their employees on how to use their new virtualization software, tools, etc.

The answer to the final question differs from business to business. Most of the time, certain departments can actually be downsized when virtualization is carried out. This is because far fewer physical machines are required. If a business has a small department, sometimes they can even eliminate this physical department and run it virtually on a single physical machine.

However, sometimes businesses may not want to downsize their departments, but instead utilize their IT department's spare physical resources. By doing this, businesses can more effectively manage and direct their physical resources to where they are required.

By carrying out virtualization, a business should be able to reduce repetitive tasks for administrators and IT professionals. This means that they can then use their knowledge more effectively for other areas of the business.

Businesses should ensure that they have a standardized set of network-wide procedures for their new virtualized environment to help to reduce the number of potential problems they may face. All employees should also be made aware of any changes to their role in the business and given necessary training if it is required. By doing this, businesses can ensure that their transition from physical to virtual is as smooth for their administrators and IT professionals, as it is for them.

## Service Management Processes

There are a number of service management processes from the ITIL framework that can play a role in virtualization. This chapter will explore some of those key processes. You will be familiar with some of these processes from the Cloud Computing Foundation program. Any new processes will be discussed in more detail.

### ***IT Financial Management***

**Objective:** To provide cost effective stewardship of the IT assets and the financial resources used in providing IT services.

IT Financial Management enables an organization to fully account for the mount spent on IT Services, and to attribute these costs to the services delivered to the organization's customers.

Using IT Financial Management to provide services with cost transparency clearly understood by the business and then rolled into the planning process for demand modeling and funding is a powerful benefit for the organization. It enables the best balance to be struck between the opportunities available for the business against the capability levels of the IT organization.

### 5.4.1 IT Financial Management and Virtualization

Once a business has moved to virtualized environment, they begin to reap the many rewards that this solution provides. Less physical machines mean less physical space is required; therefore a business's office and building costs are greatly reduced. In addition, both heat and power outputs are significantly reduced and again this reduces the running costs for a business. Networking costs are also reduced because fewer switches, hubs and wiring closets are required.

#### *Budgeting*

The Budgeting process ensures that the correct finance is available for provision of IT Services and that during the budget period that they are not over-spent. This process has a key influence on strategic and tactical plans.

It is also the means of delegating control and monitoring performance against predefined targets. It is paramount that budgets are effectively integrated within the organization and the managerial responsibility and accountability is matched and communicated in an efficient way.

As all spending affects profitability, it must be recognized that decisions about investment in IT Services and the integrated management IT Accounting discipline can help provide the competitive edge necessary for survival of an organization.

There are two types of budgeting that can be applied to virtualization:

- **Zero-based budgeting:** Starting from scratch; and
- **Incremental budgeting:** Taking last years expenditures into account.

The final budget agreed for an IT organization may include financial disciplines imposed by the organization, including:

- Limits on capital expenditure;
- Limits on operational expenditure;
- Limits on variance at any point in time, between actual and predicted spend;
- Guidelines on how the budget must be used;
- An agreed workload and set services to be delivered;
- Limits on expenditure outside the organization or group of organizations; and
- Agreements on how to cope with expectations.

## ***Information Security Management***

**Objective:** To align IT, business and information security and effectively manage in all service and IT Service Management activities.

Security objectives are met when:

- Information is available and usable when required, and the systems that provide it can appropriately resist attacks and recover from or prevent failures (availability);
- Information is observed by or disclosed to only those who have a right to know (confidentiality);
- Information is complete, accurate and protected against unauthorized modification (integrity); and
- Business transactions, as well as information exchanges between enterprises, or with partners, can be trusted (authenticity and non-repudiation).

Information Security is a system of policies and procedures designed to identify, control and protect information and any equipment used in connection with its storage, transmission and processing.

### **5.4.2 Information Security Management and Virtualization**

If a virus gains access to a physical machine, it will usually affect the entire machine. Virtual machines are separate entities, therefore any viruses or problems will be isolated within that virtual machine. Virtualization also makes it easier to revert back to a previous state; for example, an entire virtual machine can be backed up manually, at regular intervals, or using the virtual machines built in checkpoint feature.

## ***5.5 Release & Deployment Management***

**Objective:** To deliver, distribute and track one or more changes in a release into the live environment.

Often forgotten or ignored in many IT Service Management implementations or initiatives, Release and Deployment can be mistakenly seen as the poor cousin of Change Management, being of less importance and priority to both the business and IT organizations.

Much of the confusion and misunderstanding is perpetuated by the idea that Release and Deployment only focuses on the actual distribution of changes to the live environment. While timely and accurate distribution is indeed a goal of the process, the actual scope

includes all of the activities, systems and functions required to build, test and deploy a release into product and enable effective handover to service operations.

In conjunction with Change Management, Release and Deployment will enhance an organization's capabilities to develop, compile, reuse, distribute and rollback releases in accordance with defined policies that improve efficiency and reduce business disruption.

### 5.5.1 Goals and Objectives

To deploy new releases into production, transition support to service operation, and enable its effective use in order to deliver value to the customer.

Other objectives of Release and Deployment are:

- To define and agree upon Release policies, and Release and Deployment plans with customers and stakeholders;
- Ensure the integrity of constructed release packages and that they are recorded accurately in the Configuration Management System (CMS);
- Ensure that all release packages can be tracked, installed, verified, uninstalled or backed out if necessary;
- Ensure the required skills and knowledge is transferred to support staff, customers, end users, suppliers and any other relevant stakeholders; and
- There is minimal unpredicted impact on the production services, customers and service operations.

### 5.5.2 Scope

Release and Deployment works closely in conjunction with the other RCV processes to enable the quality transition of services. The role played specifically by Release and Deployment is to build, package, validate and distribute authorized service changes to the target systems or users. A *release* is a collection of authorized changes to an IT service.

### 5.5.3 Benefits

When identifying the benefits that Release and Deployment provides for, it is important to remember that it should be utilized in conjunction with the other Service Transition processes. As a result, improvements in the metrics defined for Release and Deployment may be at the expense of the other transition processes. Typical benefits seen as a result of improved Release and Deployment are:

- Delivering change, faster, at optimum cost and minimized risk;

- Assuring customers and users they can use the new or changed service in a way that supports the business goals,
- Improving consistency in implementation approach across the business change, service teams, suppliers and customers; and
- Contributing to meeting auditable requirements for traceability through Service Transition.

Well planned and implemented release and deployment will make a significant difference to an organization's service costs. A poorly designed release and deployment will, at best, force IT personnel to spend significant amounts of time troubleshooting problems and managing complexity. At worst, it can cripple the environment and degrade the live services.


### 5.5.4 Terminology

#### *Release Unit:*

A 'release unit' describes the portion of a service or IT infrastructure that is normally released together according to the organization's release policy. When defining the most appropriate release unit structure, the following factors should be taken into account:

- The ease and amount of change necessary to release and deploy the release unit;
- The amount of resources needed to build, test and distribute a release unit;
- The complexity of interfaces between the proposed unit and the rest of the services and IT infrastructure; and
- The storage and resources available in the build, test, distribution and live environments.

Based on these factors, it may mean that for critical applications the release unit is the complete set of application components, and for optional features or add-ons only the single component itself.





*Release Package:*

A release package may be a single release unit or a structured set of release units, including the associated user or support documentation that is required. Like the definition of release units, factors such as the modularity of components, the amount of change occurring and resources required will be considered when formulating a complete Release Package.

*Release Identification:*

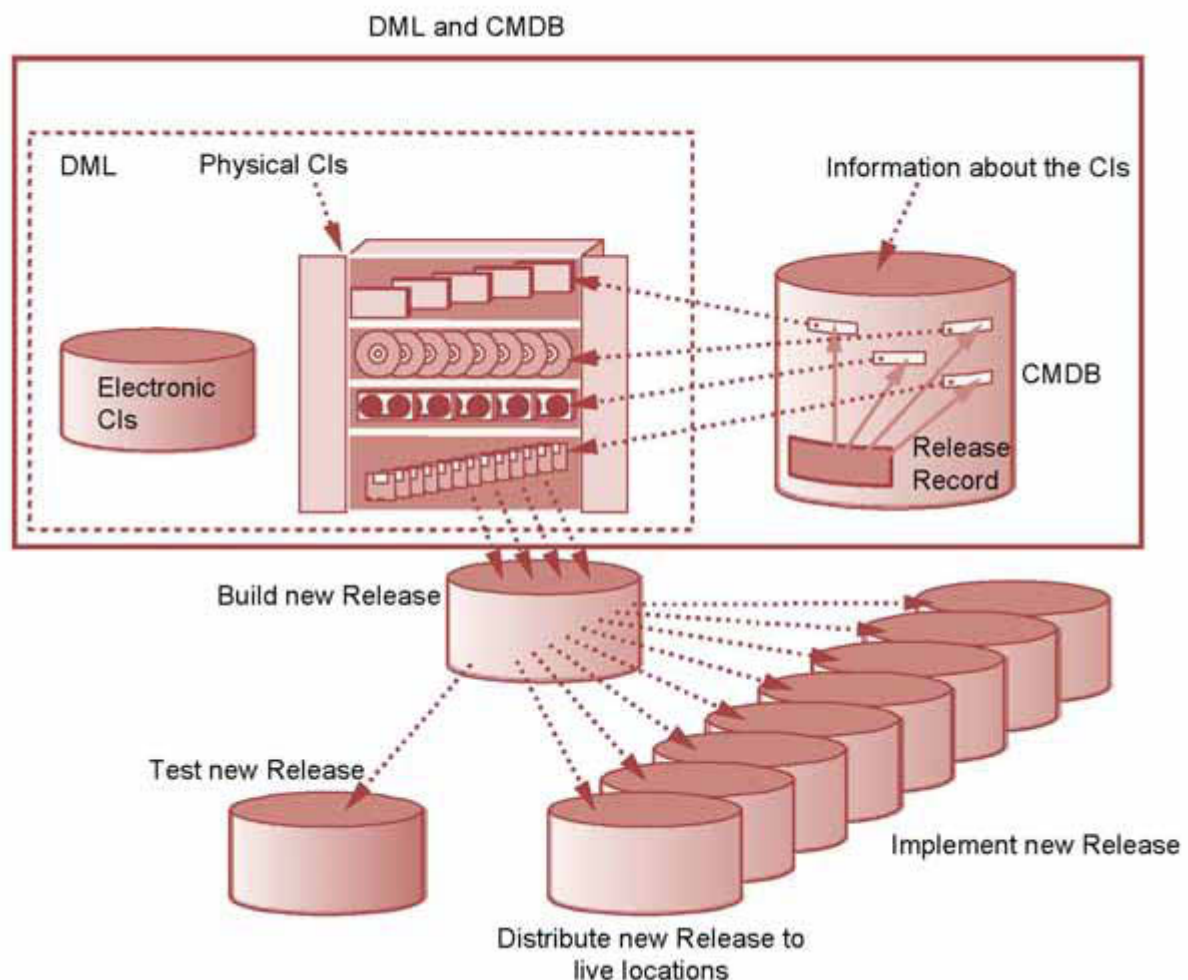
The unique release identification scheme which is normally defined in the Release Policy and conforms to any standards defined by Service Asset and Configuration Management. Typical identification schemes include the following examples:

- *Major Release:* Banking\_System v1, v2, v3 etc
  - Major roll-out of new hardware and/or software.
- *Minor Release:* Banking\_System v1.1, v1.2, v1.3 etc
  - A few minor improvements and fixes to Known Errors.
- *Emergency Fix:* Banking\_System v1.1.1, v1.1.2 etc
  - A temporary or permanent Quick Fix for a Problem or Known Error.

*Definitive Media Library (DML):*

DML is the secure library where the definitive authorized versions of all media CIs (typically software) are stored and protected. The DML should include definitive copies of purchased software (along with license documents and any controlled information) as well as software developed on site.

The DML includes both electronic and physical copies of definitive software releases, with Release and Deployment being responsible for any additions, modifications, removals or maintenance that needs to be performed.

*Definitive Spares:*

Physical storage of all **spare IT components and assemblies maintained at the same level as within the live environment**. New IT assemblies are stored here until ready for use, and additional components can be used when needed for additional systems or in the recovery from Incidents. Like the DML, details of these assemblies are recorded in the

CMDB, and their management and deployment the responsibility of the Release and Deployment process.

#### *Build Management:*

The software, hardware and documentation that comprise a release unit should be assembled in a controlled manner to ensure a repeatable process. This should include automation where possible for its compilation and distribution, which for large organizations can significantly reduce the Total Cost of Ownership (TCO) for the services involved.

### **5.5.5 Triggers and Interfaces**

The primary interfaces of release and deployment exist with Change Management and the surrounding Service Transition processes. Other inputs will also be provided from Service Strategy and Service Design to ensure that the requirements for value provision have been met.

The inputs to release and deployment include:

- Authorized RFCs;
- Service Packages;
- Service Design Package;
- Service Acceptance Criteria;
- Service Management policies and standards (including the Change Policy);
- Build Models and plans; and
- Exit and entry criteria for each stage of release and deployment.

The outputs include:

- Release and deployment plan;
- Updated RFCs for any required activities;
- Updated service catalogue reflecting any service changes;
- New or modified services;
- New or modified processes;
- Skilled and knowledgeable support staff;
- End users with capabilities to use the service;
- SLAs, OLAs, UCs;
- Deployment plans and packages; and
- Service Transition Reports.

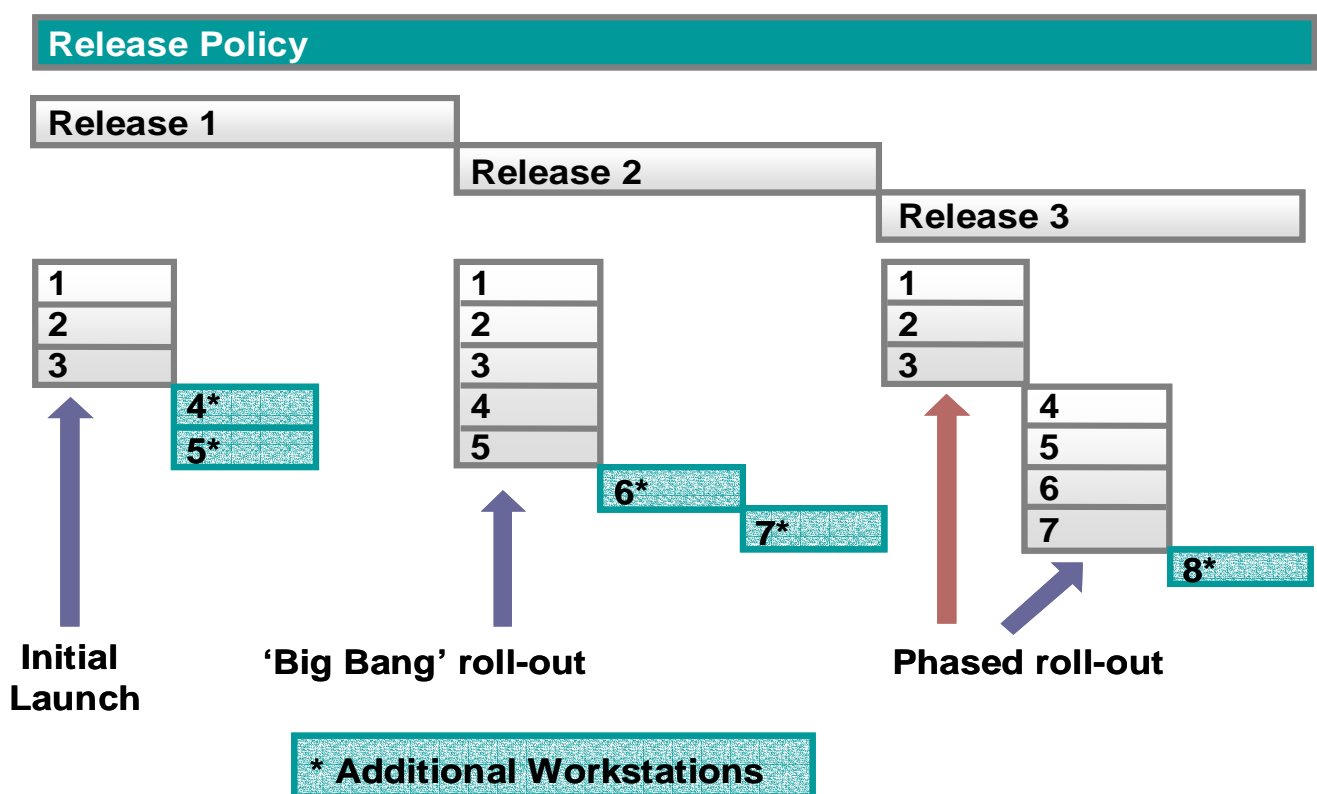
### 5.5.6 Release Design Options and Considerations

When planning individual releases or defining the policies that should exist, consideration about the potential impact and need for resources will affect how releases will be deployed to the target locations. The common options for deploying releases are described below:

#### *Big Bang or Phased Approach*

**Big Bang:** Where the new or changed service is deployed to all user areas in one operation. This will often be used when introducing an application change and consistency of service across the organization is considered important.

**Phased Approach:** The service is deployed to a part of the user base initially, and then this operation is repeated for subsequent parts of the user base via a scheduled rollout plan. This will be the case in many scenarios such as in retail organizations for new services being introduced into the stores' environment in manageable phases.



### *Push or Pull*

**The Push Approach:** Service component is deployed from the centre and pushed out to the target locations. In terms of service deployment, delivering updated service components to all users, either in big bang or phased forms is the use of a push approach, since the new or changed service is delivered into the users' environment at a time not of their choosing.

**The Pull Approach:** Used for software releases. Software is made available in a central location but users are free to pull the software down to their own location at a time of their choosing, or when a workstation restarts. The use of 'Pull' updating a release over the internet has made this concept significantly more pervasive. A good example is virus signature updates, which are typically pulled down to update P.C's and servers when it best suits the customer; however at times of extreme virus risk this may be overridden by a release that is pushed to all known users.

Pull approaches do not rest so heavily on accurate configuration data and they can trigger an update to user records. This may be through new users appearing and requesting downloads or expected users not doing so, triggering investigation into their continued existence.

### *Automated or Manual*

**Automated:** Helps to ensure repeatability and consistency. The time required to provide a well-designed and efficient automated mechanism may not always be available or viable. Typical examples of activities capable of a degree of automation:

- Discovery tools aid release planning;
- Automated builds reduce time taken - this in turn can resolve scheduling conflicts and delays; and
- Automated configuration baselines procedures save time and reduce errors in identifying the status of CI's and releases during build, test and deployment etc.

**Manual:** Important to monitor and measure the impact of many repeated manual activities as they are likely to be inefficient and error prone. This will ultimately slow down the release team and create resource and capacity issues that affect the agreed service levels.

### 5.5.7 Release Policy

A Release Policy is the formal documentation of the overarching strategy for releases and was derived from the Service Design phase of the Service Lifecycle. It is the governing policy document for the process and must accommodate the majority of releases being implemented. Typical contents of a Release Policy include:

- Level of infrastructure to be controlled by releases;
- Preferred structure and schedules for Release Packages;
- Definition of major and minor releases, emergency fixes;
- Expected deliverables for each type of release;
- Policy on the production and execution of back out plans;
- How and where releases should be documented;
- Blackout windows for releases based on business or IT requirements;
- Roles and responsibilities defined for the Release and Deployment process; and
- Supplier contacts and escalation points.

### 5.5.8 Release and Deployment Activities



#### Key Points:

The Release Policy is the overarching strategy for Releases and was derived from the Service Design phase of the Service Lifecycle.

The Release Plan is the operational implementation for each release.

The Deployment Plan is the documented approach for distributing a single Release.

**Overview of Important Steps:**

1. Release planning
2. Preparation for build, test and deployment
3. Build and test
4. Service test and pilots
5. Plan and prepare for deployment
6. Perform transfer. Deployment and retirement
7. Verify deployment
8. Early Life Support
9. Review and close the deployment
10. Review and close Service Transition

**1. Release Planning**

Any plans created for the release and deployment will need to be integrated with the overall Service transition plan, and conform to any policies that have been defined. For each release, plans should be authorized by Change Management and used to assist in the evaluation of the risk, impact and resource requirements for components of the change. Typically the release and deployment plans should document the:

- Scope and content of the release;
- The risk assessment for the release;
- Affected stakeholders;
- Teams responsible for the release; and
- Communication strategy to be used during the release and deployment process.

Plans should take into account the acceptance criteria that exist for the release and when authorization points will verify a pass or fail. The processes of Evaluation and Service Validation and Testing will be integrated here to assist in the determination whether to continue, pause or revert to previous stages of the release.

*Build and test planning*

The approach taken for building, testing and maintaining the controlled environments to production will need to be planned in order to enable optimum use of resources for the development of the release. The activities that occur here are:

- Developing build plans based on the Service Design Package and defining any environment requirements;
- Scheduling the resources and time required to setup the environments;
- Testing the build and compilation procedures;
- Scheduling the build and compilation activities;
- Assigning resources, roles and responsibilities for any key activities; and
- Defining the build exit and entry criteria.

Environments that may be utilized during this period include:

- Build environments;
- Testing and integration environments;
- Deployment environments;
- Pilot environments; and
- Backup and recovery environments.

### *Utilizing Pilots*

Pilots may be useful for testing the service with a group of participants that are representative of the broader end-user community. For this to be effective the scope needs to be carefully determined, as being either too large or too small will likely result in some negative impact to the overall success and quality of the release and deployment process.

Pilots should include mechanisms by which feedback can be gathered about various aspects of the release and associated processes. For complex releases or in large and diverse organizations it may be appropriate to use more than one pilot for address the different implementation and support issues that exist.

### *Deployment Planning*

There are many factors that will be considered when choosing the most appropriate deployment strategy (refer back to deployment options). Questions that must be answered include:

1. What needs to be deployed?
2. Who are the users?
3. Are there location dependences?
4. Where are the users?
5. Who else needs to be prepared well in advance? (training, security etc)
6. When does the deployment need to be completed?
7. Why is the deployment happening?
8. What are the critical success factors and exit criteria?
9. What is the current capability of the service provider?

### *Financial/Commercial Planning*

Where necessary, various financial and commercial factors will need to be assessed before the deployment activities, including:

- Working capital;
- Contracts and licenses;



- Funding; and
- Intellectual property requirements.

## 2. Preparation for build, test and deployment

Before the actual building of the release occurs, the release design must be validated against the requirements defined for the new or changed service offering. This should be an independent evaluation that checks the release will deliver the predicted outcomes, and any issues documented in an interim evaluation report.

### *Training of involved release and deployment staff*

In many cases the introduction of a release may require additional training for the release, deployment, build and test teams. Such training could be related to the:

- Service management processes to be used;
- Changes in security or health and safety procedures;
- Understanding of the Service Design documentation and plans; and
- Technology being utilized for the release.

## 3. Build and test

Wherever possible, repeatable practices and reusable components should be utilized to during the build and test of releases. This includes managing the:

- Build, test and packaging environments;
- Compilation and packaging tools; and
- Configuration of the releases themselves:
  - Version control;
  - Documentation templates for testing and validation; and
  - Access rights and security procedures.

### *Release and build documentation*

Documentation templates, procedures, knowledge bases and other guidance should be consistently available to support the release team in the activities performed. Typical documentation that will be used by the release teams include:

- Contract agreements;
- Purchase requests;
- Health and Safety guidelines;
- Security policies;
- Licence agreements;
- Procedures for:
  - Distributing software;

- Delivering, moving and installing equipment;
- Wiping sensitive data and media; and
- Publishing, sharing and archiving knowledge, information and data.

#### *Acquire and test required components*

Release and Deployment should be interfaced with the organization's existing procurement processes to acquire any required components for the release. This will save time and effort in verifying assets, capturing and recording information, ensuring proof of licence and triggering updates to the Asset Management System.

As part of the overall Service Validation and Testing, each of the individual components should be tested to verify that any quality criteria has been met, initiating action where it's not met.

#### *Release Packaging*

Build management procedures, tools and checklists should be utilized as part of the release packaging, to provide repeatable practices and expected outcomes. When a definitive package has been assembled, a baseline should be taken of the release package and the correct versioning and naming conventions applied.

#### *Managing the build and test environments*

The need for multiple environments in which to build and test will depend on the size, complexity, frequency and impact of the releases being managed. Test environments should be protected using a range of testing best practices (see also Service Validation and Testing), and appropriate access to these environments based on the priorities defined.

Automating the installation of systems and software reduces the workload of people, but also requires testing of the scripts and mechanisms that will be used.

### **4. Service testing and pilots (See *Service Validation and Testing*)**

As part of a coordinated effort with Service Validation and Testing, testing and validation must be performed at multiple levels. With particular focus on the release itself, *service rehearsals* may be used, which simulate as much of the service as possible in an extensive and widely involved practice session. This would normally occur after other pilots have run, and is designed to be the last measure to detect any potential issues that will arrive during or after the deployment to the live environment.

### *Pilots*

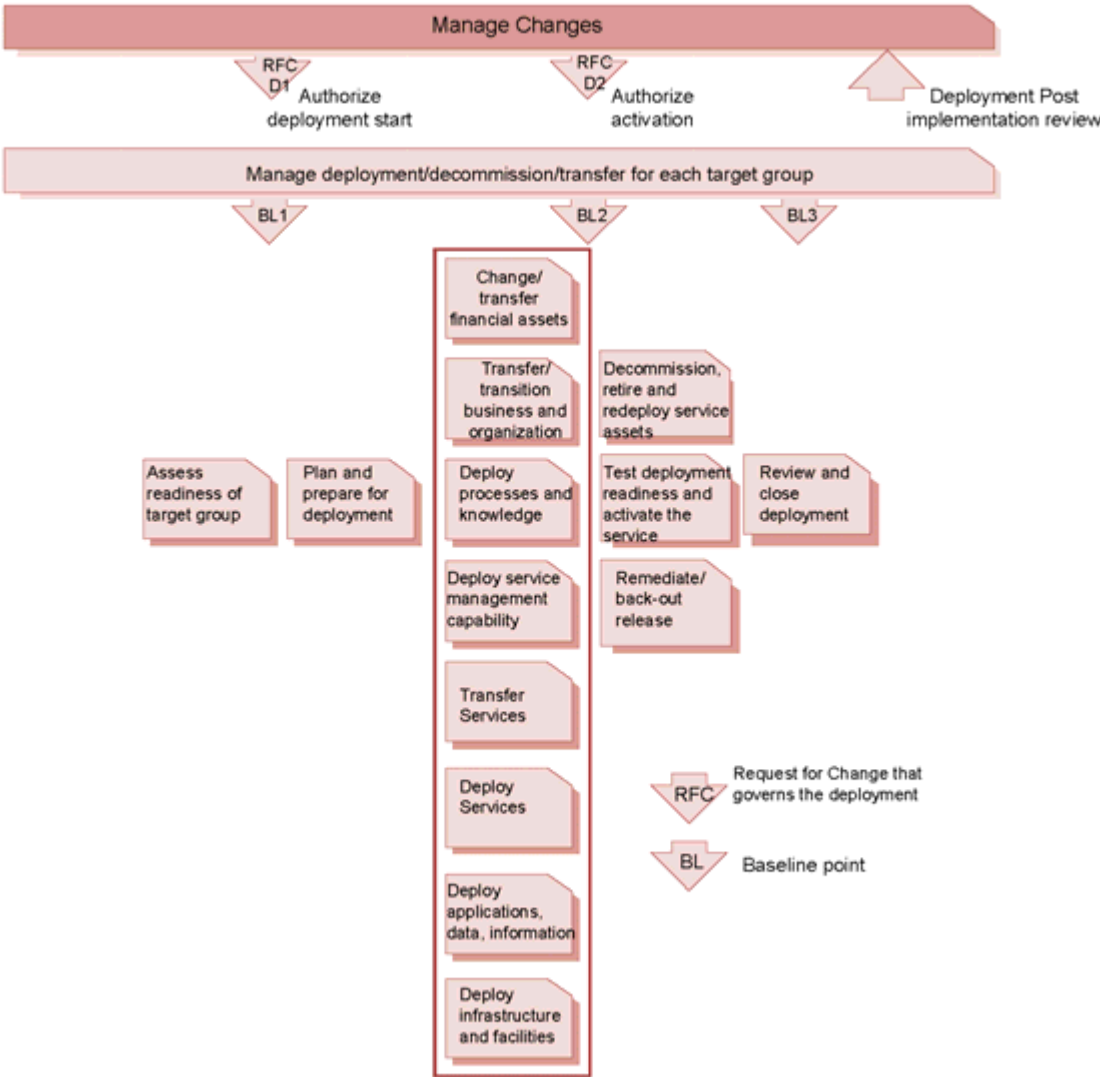
Previous planning should have already identified what pilots will be used as part of the overall release and deployment. Key actions to take during pilots are to:

- Train any people involved;
- Document any required procedures;
- Continue communication and engagement with customers and users;
- Determine the levels of support required for the actual release;
- Discover and fix issues wherever possible before the final deployment; and
- Document improvements where appropriate and incorporate them into future plans.

## **5. Plan and prepare for deployment**

At this stage the focus is on preparing the organization and people for organizational change and to refine deployment plans that have been documented. These plans should include guidance regarding:

- Risk mitigation plans;
- Disposal and retirement plans;
- The logistics for delivery;
- Knowledge transfer;
- Mobilizing users to be ready to use the service; and
- Mobilizing the support staff for service readiness.



## **6: Perform transfer, deployment and retirement**

During the actual implementation the activities performed can be grouped under the following tasks:

1. Transfer financial assets;
2. Transfer changes required to business/organization;
3. Deploy processes and materials;
4. Deploy Service Management Capability;
5. Transfer service;
6. Deploy service;
7. Decommissioning and service retirement; and
8. Remove redundant assets.

These activities will need to be modified to accommodate any items specified in the deployment plan as part of the acceptance criteria for 'go live'.

## **7: Verify deployment**

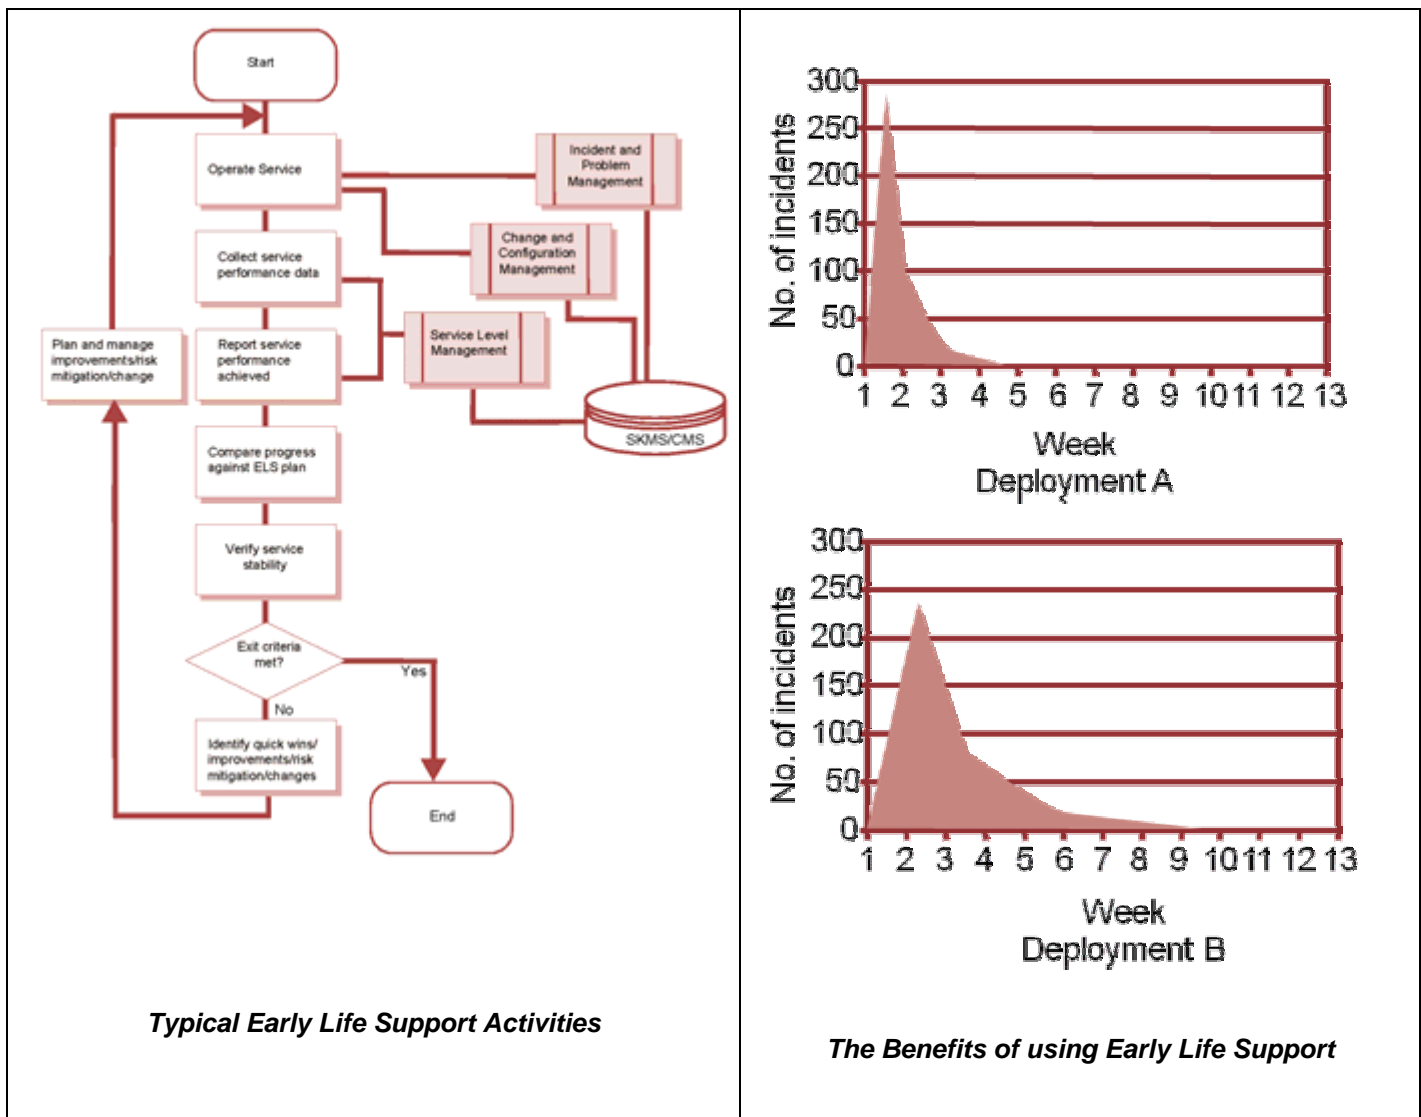
Once the activities for the deployment of releases are finished, verify that users are capable of operating the service. Verification should ensure that:

- The service/release components are in place by means of a configuration audit;
- Documentation has been updated accordingly;
- Roles and responsibilities have been correctly assigned; and
- The measurement and reporting systems are established to measure performance of the service.

Any noted deviations from plans or other items raised should be documented in order to improve future implementations and processes used.

## **8: Early Life Support**

The quality of transition to Service Operation is a crucial element to the success of the overall service change that is being implemented. Rather than simply palm off support post-deployment, the release and deployment teams should assist in managing any calls, incidents and problems that are detected in the early of the new or modified service. This enables more stability in this vulnerable period, increased customer and user satisfaction, enhanced learning and better momentum for continual improvement. The resource allocation from these teams will then be gradually reduced while Service Operation takes on more responsibility for support.



The example shown in the figures above demonstrates how the number of incidents for deployment A was significantly reduced through the use of Early Life Support, including the training of users and staff, and the transfer of knowledge to service desk staff (including knowledge-base articles).

The acceptance criteria for finalizing Early Life Support should be agreed early in the process, which might include such conditions as:

- Users can use the service effectively and efficiently for their business activities;
- Service delivery is managed and controlled across any service provider interfaces;
- Service levels and performance standards are being consistently achieved;
- Stakeholder groups verify the transfer of required knowledge; and
- All deliverables required for the release have been signed off.

## 9: Review and close the deployment

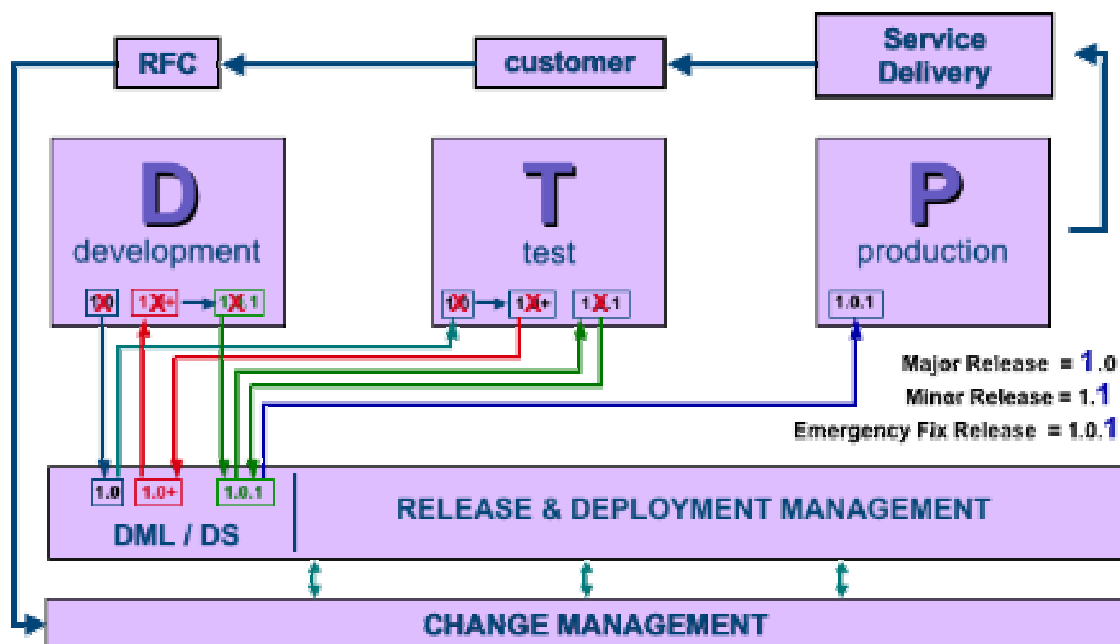
A formal review of deployments should be performed for all releases of a determined level of size, impact, cost or risk to the organization. The review seeks to ensure that all requirements for the release have been met and to identify any potential improvements that can be made. Items that should be reviewed include:

- Any quality criteria deviations that exist;
- Any open actions or necessary fixes that have been identified;
- Review open changes;
- Review performance targets and achievements;
- Experiences and feedback from customers, users and staff involved with the deployment;
- All problems and known errors are documented and accepted by the business and/or suppliers; and
- Check that any redundant assets (including licences) have been removed.

## 10: Review and close Service Transition

The final step required in finalizing the completion of Service Transition is a formal review appropriate to the relative scale of the change. This will use the process of Evaluation and is driven by Change Management, which will verify successful completion and that the handover to Service Operation is complete.

The “lessons learnt” should be documented to provide any improvement opportunities that can be made and developed by Continual Service Improvement for future transitions.



## Knowledge Management

**Objective:** To enable organizations to improve the quality of management decision making by ensuring that reliable and secure information and data is available throughout the service lifecycle.

The primary purpose of Knowledge Management is to improve efficiency by reducing the need to rediscover knowledge. This requires accessible, quality and relevant data and information to be available to staff.

The goals of Knowledge Management are to:

- Enable the service provider to be more efficient and improve quality of service, increase satisfaction and reduce the cost of service;
- Ensure staff have a clear and common understanding of the value that their services provide to customers, and the ways on which benefits are realized for the use of those services; and
- Ensure that, at a given time and location, service provider staff have adequate information on:
  - Who is currently using their services;
  - The current states of consumption;
  - Service delivery constraints; and
  - Difficulties faced by the customer in fully realizing the benefits expected from the service.

The quality of decision making within the Service Lifecycle depends on the ability and understanding of those parties involved, the understanding of the benefits and consequences of actions taken, and the analysis of any of the surrounding issues involved. All of this in turn depends on the availability of accurate and timely knowledge, information and data provided in a way that can be easily accessed and interpreted by the appropriate parties.

Knowledge rests on the management of the information and data that underpins it. To be efficient this process requires an understanding and maintenance of the following key activities:


- Establishing data and information requirements;
- Define the information architecture; and
- Establishing data and information management procedures.



### 5.5.9 Policies and principles of Knowledge Management

#### *Data, Information, Knowledge, Wisdom (DIKW)*

Knowledge Management is usually seen within the context of the DIKW structure seen below.



#### **Data:**

Data is a set of discrete facts. Most organizations capture significant amounts of data every day. Knowledge Management activities seek to improve the capabilities for capturing, analyzing, synthesizing data and transforming it into information.

#### **Information:**

Information comes from providing context to data. This usually requires capturing various sources of data and applying some meaning or relevance to the set of facts. Knowledge Management focuses on measures for capturing, finding and reusing information so that work is not duplicated.

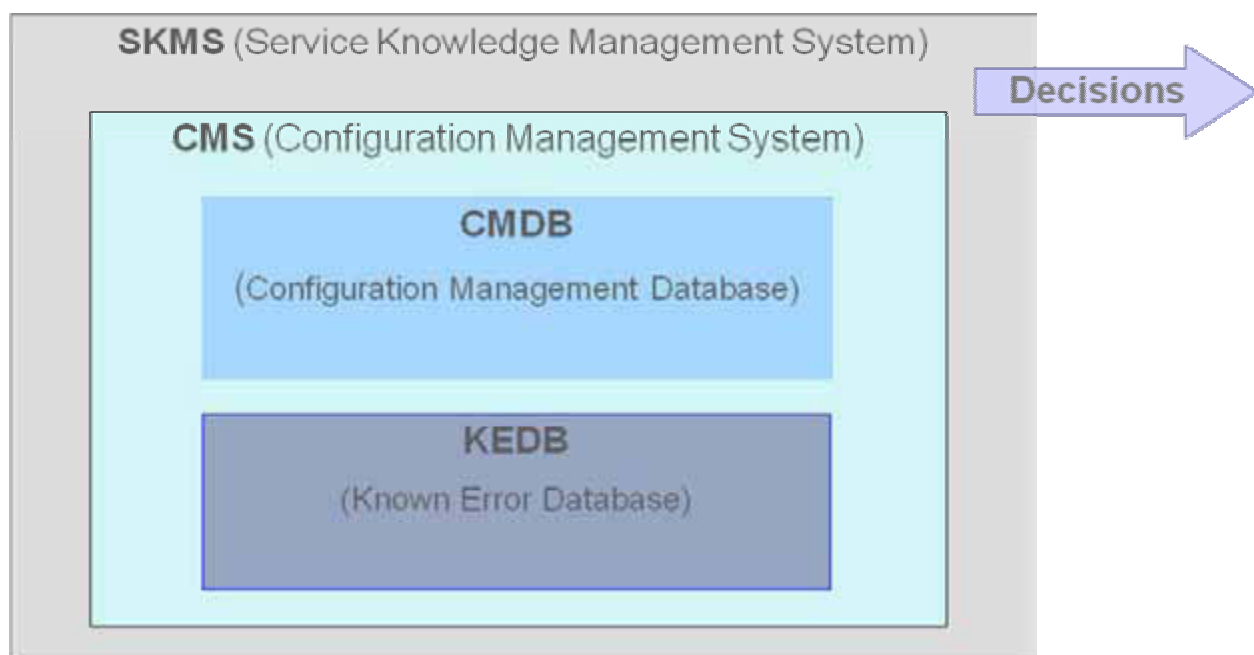
**Knowledge:**

Knowledge is composed of the experiences, ideas, insights and judgments from individuals and from their peers. This usually requires the analysis of information, and is applied in such a way to facilitate decision making.

**Wisdom:**

The ability to ultimately discern material and having the application and contextual awareness to provide a strong common sense judgment. The use of wisdom ultimately enables an organization to direct its strategy and growth in competitive market spaces.

We can use tools and databases to capture data, information and knowledge, but wisdom cannot be captured this way, as wisdom is a concept relating to abilities to use knowledge to make correct judgments and decisions.

**5.5.10 The Service Knowledge Management System (SKMS)****SKMS: Service Knowledge Management System (SKMS):**

The SKMS describes the complete set of tools and databases that are used to manage knowledge and information, including the Configuration Management System as well as other tools and databases. The SKMS stores, manages, updates and presents all information that an IT service provider needs to manage the full lifecycle of its services. The main purpose of the SKMS is to provide quality information so that informed decisions can be made by the IT service provider.

Whereas the CMS focuses on providing information relating to the configuration of the IT infrastructure, the SKMS has a broader scope (as implied by the diagram) which includes anything pertaining to the needs of service management, including:

- Experience of staff;
- Records of peripherals;
- Supplier & Partner requirements and abilities; and
- Typical and anticipated user skill levels.

## ***Incident Management***

**Objective:** To restore normal service operation as quickly as possible and minimize the adverse impact on business operations, thus ensuring that the best possible levels of service quality and availability are maintained.

Incident Management has developed over time to become one of the most visible and mature ITIL processes for any organization, largely driven by the need to reduce the business impact of disruptions to IT services. While any effective implementation does balance the efforts towards the various phases of the Service Lifecycle, as Incident Management can be easily demonstrated to have a business benefit, it typically receives more attention and funding than other areas of service management. This section will explain the activities and techniques that represent best practices for Incident Management.

### **5.5.11 Incident Management and Virtualization**

When a business implements virtualization, there are many things that can easily be overlooked or forgotten about; for example, identifying impacts to Incident Management and Technical support. When a business runs software or applications in a normal physical environment, technical support will often automatically be provided as a part of their purchase.

However, with the advent of virtualization technology, things are not as straightforward as they once were. For example, if a vendor's application is run in a virtualized environment, where it was not originally intended to, the vendor may not be contractually required to continue to offer technical support. This is an important consideration for the business.

Traditionally, software applications ran within a single environment. This meant that technical support staff could all follow a straightforward script when troubleshooting. This was until the client/server model became well established. The client/server model allowed software to be distributed and run from many different locations. This has complicated the

role of technical support staff, as there were now multiple scenarios to be looked at when troubleshooting and a single problem could have been caused by a number of different issues.

Virtualization further adds to this complexity, the more programs or software applications that are running on a single physical machine, the more potential ways there are for a problem to occur. This obviously makes the job for technical support staff much more challenging.

Further challenges can be faced with a lack of appropriate testing for application software and/or a lack of knowledge and understanding of staff, within the virtualization environments.

As most software vendors design their software to be used in a physical environment, the testing environment will also be based on a physical environment as opposed to a virtualized one. Therefore, most vendors will be unable to confirm whether or not their software can run stably in a virtualized environment or not.

This also means that a software vendor's technical support staff would be unlikely to have the knowledge to offer support to a business that uses virtualization technology.

As an example, a business has a critical application stop working after migrating to a virtualized environment. They contact their technical support staff who manage that particular application, to find that technical support does not offer support, or does not have the knowledge to offer support, for that businesses virtualized environment.

## **5.6 *Change Management***

**Objective:** To ensure all changes are assessed, approved, implemented and reviewed in a controlled manner.

The ability to control and manage changes to defined IT services and their supporting elements is viewed as fundamental to quality service management. When reviewing the typical strategic objectives defined for an IT service provider, most of these are underpinned by the requirement of effective change control. These include strategies focusing on time-to-market, increased market share or high availability and security platforms, all of which require a controlled process by which to assess, control and manage changes with varying levels of rigour.

Changes arise for a number of reasons, including:

- Requests of the business or customers, seeking to improve services, reduce costs or increasing ease and effectiveness of delivery and support; and
- From internal IT groups looking to proactively improve services or to resolve errors and correct service disruption.

The process of Change Management typically exists in order to:

- Optimize risk exposure (defined from both business and IT perspectives);
- Minimize the severity of any impact and disruption; and
- Deliver successful changes at the first attempt.

To deliver these benefits it is important to consider the diverse types of changes that will be assessed and how a balance can be maintained in regards to the varying needs and potential impacts of changes. In light of this, it is important to interpret the following Change Management guidance with the understanding that is intended to be scaled to suit the organization and the size, complexity and risk of changes being assessed.

### 5.6.1 Goals and Objectives

To ensure that **standardized methods and procedures** are used for controlled, efficient and prompt handling of all changes, in order to **minimize the impact** of change-related Incidents upon service quality, and consequently to improve the day-to-day operations of the organization.

***“Remember: Not every change is an improvement, but every improvement is a change!”***

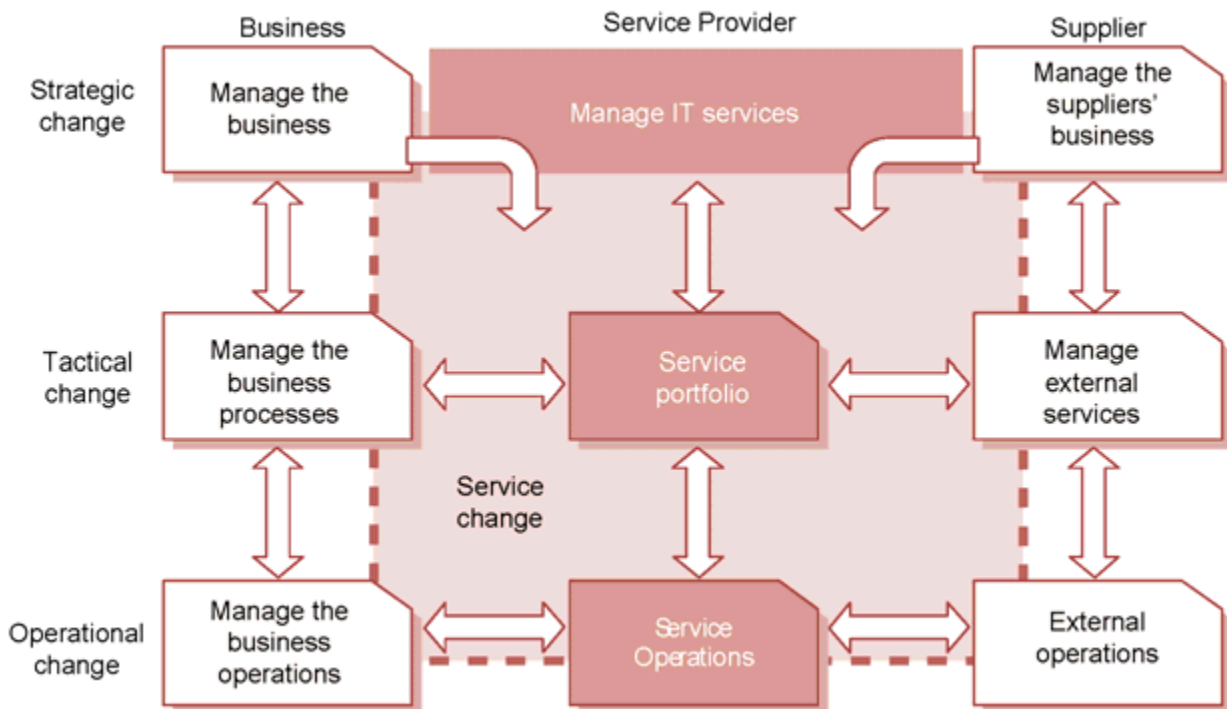
Change Management’s purpose is also to ensure that:

1. All Changes to service assets and configuration items (CIs) are recorded in the Configuration Management Systems (CMS).
2. Overall business risk is optimized.

### 5.6.2 Scope

The term Change is often defined in varying ways, however the best definition of a service change is: *“Any alteration in the state of a Configuration Item (CI). This includes the addition, modification or removal of approved, supported or baselined hardware, network, software, application, environment, system, desktop build or associated documentation.”*

It is important, however, that every organization defines those changes which lie outside the scope of their service change process (such as operational or business process and policy changes).



The figure above demonstrates the typical scope of the Change Management process for an IT Service Provider and how it interfaces with the business and suppliers at strategic, tactical and operational levels. As discussed in 4.1.4, Service Portfolios provide the clear definition of all planned, current and retired services.

### 5.6.3 Designing and Planning

It is generally advised that the Change Management process should be planned in conjunction with Release & Deployment, and Service Asset & Configuration Management. These processes together will help in the evaluation of impact, needs, timings and overall risk for changes being assessed.

The checklist of typical requirements when designing the Change Management process includes:

- Regulatory, policy or other compliance requirements;
- Documentation requirements;
- Identification of impact, urgency and priority codes for changes;
- Roles and responsibilities involved;
- Procedures required;
- Interfaces to other Service Management processes (e.g. Problem Management);
- Toolset requirements to support Change Management; and
- Configuration Management interfaces.

### 5.6.4 Change Management Policies

There must be policies and standards defined that clarify, for internal and external providers, who must do what, when and what the consequences of non-compliance to the policy will be. These policies are typically created with information relating to:

- Creating a culture of Change Management – with zero tolerance on unauthorized changes;
- Aligning the service Change Management process with business, project and stakeholder Change Management;
- Prioritization of changes;
- Establishing accountability and responsibilities;
- Segregation of duty controls;
- Establishing a single focal point for changes;
- Prevent people who are not authorized to make changes from accessing production environment;
- Integration with other Service Management processes;
- Changing windows;
- Performance and risk evaluation; and
- Performance measures for the process.

### 5.6.5 Change Models

The definition of different process models will allow an organization to maintain a balance between providing an appropriate level of control for changes without causing bottlenecks or restricting business growth. Change Models define how various categories of changes are assessed & authorized, with different mechanisms and activities used to process and deliver changes based on the change type. The defined Change Models should also include:

- What steps should be taken to manage the change;
- Roles and Responsibilities;
- Timescales and thresholds for actions; and
- Escalation procedures.

Change Models defined within ITIL include the following:

**NORMAL  
Change:**

A change that follows all of the steps of the change process. It is assessed by either a Change Manager or Change Advisory Board. Normal Changes will often be further defined by the relative impact and complexity, which will escalate the change for

assessment to the most appropriate person or group.

See Figure 6.3 for the example process flow for normal changes.

**STANDARD  
Change:**

A *pre-approved* change that is low risk, relatively common and follows a procedure or work instruction. E.g. password reset or provision of standard equipment to a new employee. RFCs are not required to implement a Standard Change, and they are logged and tracked using a different mechanism, such as a **service request**. While standard changes are effectively pre-approved by Change Management, they may still require forms of authorization such as other groups such Human Resources (HR) or Financial departments.

The main elements of a standard change are that:

- Authority is effectively given in advance;
- The tasks are well known, documented and proven;
- There is a defined trigger to initiate the Request For Change (RFC);
- Budgetary approval is typically defined or controlled by the requester; and
- The risk is usually low and always well understood.

Over time and as the IT organization matures, the list of standard changes should increase in order to maintain optimum levels of efficiency and effectiveness.

**EMERGENCY  
Change:**

A change that must be introduced as soon as possible.  
E.g. to resolve a major incident or implement a security patch.

The change management process will normally have a specific procedure for handling Emergency Changes quickly, without sacrificing normal management controls. Organizations should be careful to ensure that the number of emergency changes be kept to a minimum, because they are typically more disruptive and prone to failure.

To enable this to occur, methods of assessment and documentation are typically modified; with some documentation occurring after the change has occurred.



### 5.6.6 Triggers and Interfaces

Requests for change can come from anywhere within service lifecycle or from interfaces with the business, suppliers or other stakeholders. The inputs required for the change and outputs produced will vary depending on the type defined and whether it is strategic, tactical or operational in nature.

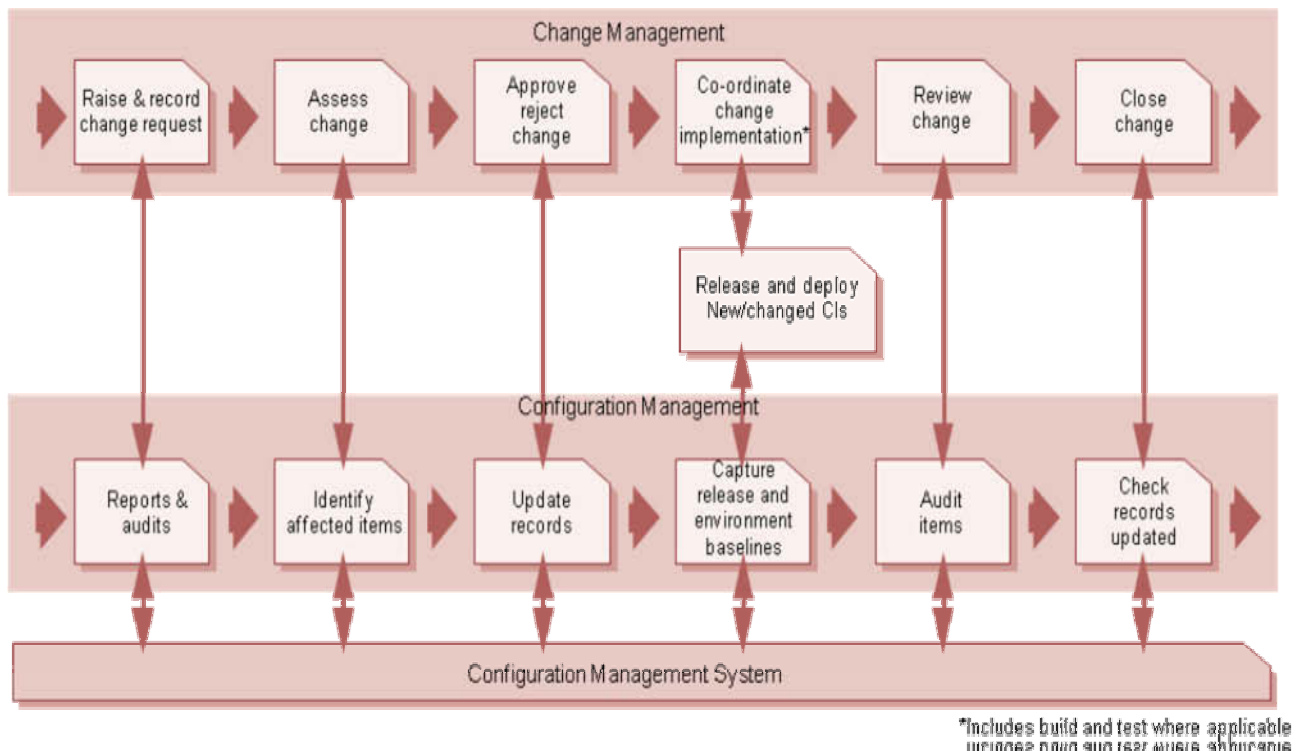
Normal **inputs** include:

- Request for Changes (RFC);
- Policy and strategies for change and release;
- Plans (change, transition, release, deployment, test, evaluation and remediation);
- Current change schedule and projected service outage (PSO);
- Current or affected assets and configuration items; and
- Test results, test reports and evaluation reports.

Normal **outputs** from the process will be:

- Rejected RFCs;
- Approved RFCs;
- Changes to services, infrastructure components resulting from approved RFCs;
- Updated change schedule;
- Revised PSO;
- Authorized change plans; and
- Change decisions, documents, records and reports.

Interfaces with other service lifecycle processes should also be considered for a coordinated approach to quality service management. Additionally the integration of Change Management with Business Change processes, project management and supplier management is recommended to enable seamless level of IT quality to be delivered.



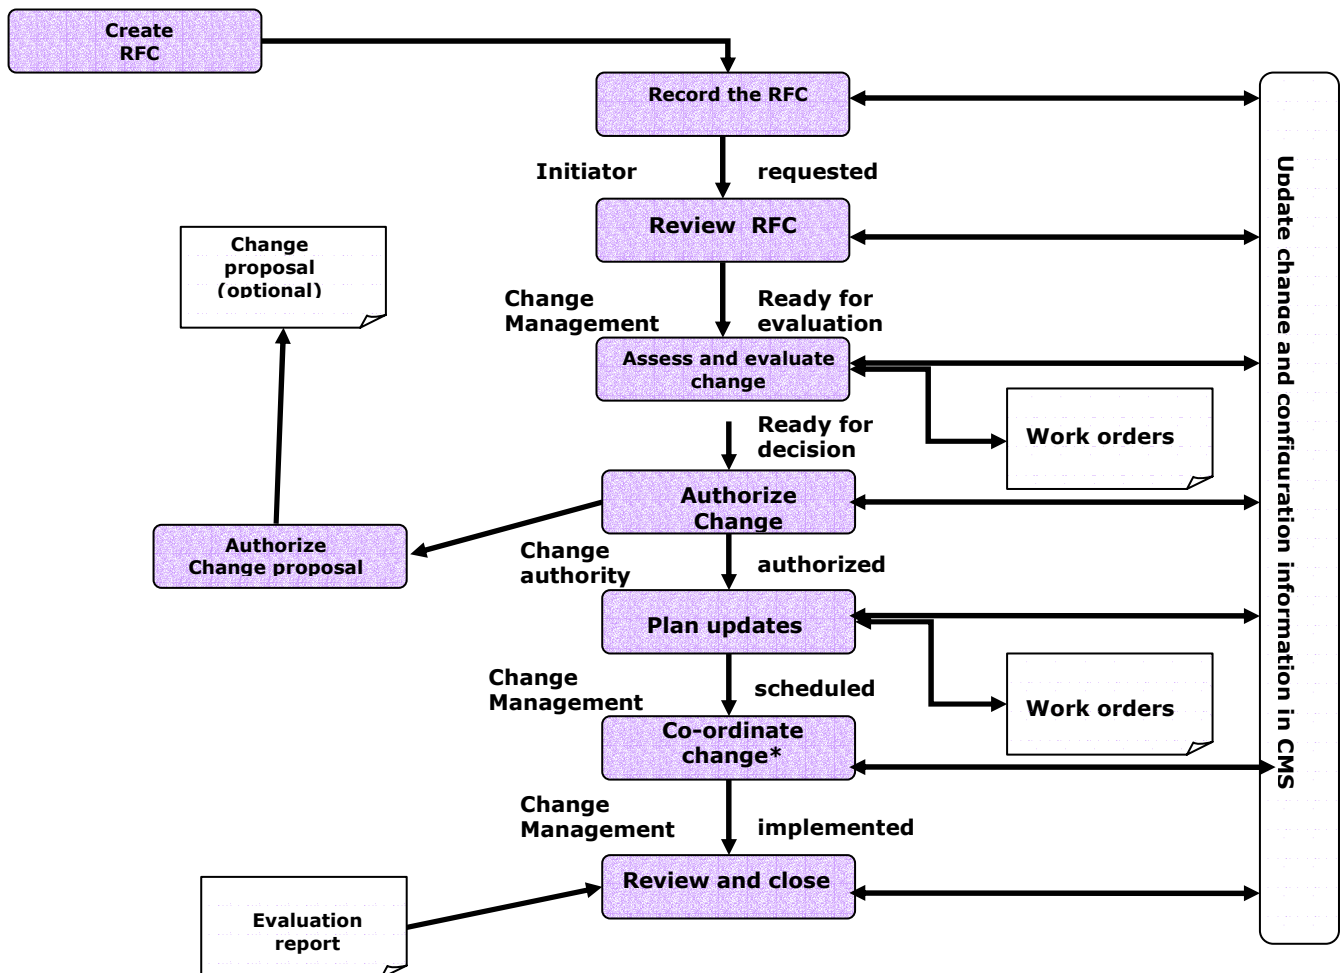
There can be various types of Change Requests to initiate Change Management processes, including:

- RFCs;
- Service Desk calls; and
- Project Initiation Documents.

Different types of change require different types of change request. An organization needs to ensure that appropriate procedures and forms are available to cover the anticipated requests from the different parties and stakeholders involved. Avoiding a bureaucratic approach to documenting a minor change alleviates some of the cultural barriers and resistance that may exist when adopting the Change Management process.

### 5.6.7 Change Management Activities

The following diagram represents the typical activities involved for normal changes that have been identified. The actual steps and procedures need to be further refined depending on any specific Change Models that have been created.



#### Overview of Important Steps:

1. The RFC is recorded.
2. Review the RFC for classification and prioritization.
3. Assess and evaluate the change – may require involvement of CAB or ECAB.
4. Authorization or rejection of the change.
5. The change is scheduled.
6. Work orders are issued for the build of the change (but carried out by other groups).
7. Change Management coordinates the work performed.
8. The change is reviewed.
9. The change is closed.

### **1. The RFC is recorded.**

The change is raised by a request from the initiator. The level of information recorded for a change depends largely on the size and impact of the change. Some information is recorded initially and some information updated as the change document progresses through its lifecycle. This may be recorded directly on the RFC form and details of the change and actions may be recorded in other documents and referenced from the RFC such as business cases.

For a major change with significant organizational and/or financial applications, a change proposal may be required, which will contain a full description of the change together with a business and financial justification for the proposed change. The change proposal will include sign off by appropriate levels of business management.

### **2. Review the RFC for classification and prioritization.**

To ensure Change Management is using an appropriate level of control based on factors of risk, cost and complexity, an initial review should act as a filtering mechanism to apply the correct Change Model (classification), identify the relative priority of the change, and to ensure that the required details are supplied. Procedures should stipulate that, as changes are logged, Change Management reviews each change request and return any that are:

- Totally impractical;
- Repeats of earlier RFCs; or
- Incomplete submissions.

These requests will be returned to the initiator, together with brief details of the reason for the rejection, and the log should record this fact. There should be an opportunity to appeal, via normal management channels, and should be incorporated within the procedures.

### **3. Assess and evaluate the change.**

All changes will be assessed for their relative potential impact, risk and resource requirements. Depending on the Change Model that has been applied, this assessment may require involvement from:

- The Change Manager and Change Owner for local authorization;
- The Change Advisory Board (representing all key stakeholders);
- The IT Management (Steering) Board; or
- Business Executive Board.

The scope of potential impacts on services for failed changes is wide, so the assessment needs to identify potential for:

- Impact on customer's business operation;
- Effect on SLAs, baselines, service models, security etc;
- Impact on other services;
- Impact on non-IT infrastructures;
- Effect of not implementing;
- Cost and staffing implications;
- Current Change Schedule;
- Ongoing resources required after implementation; and
- Impact on continuity plan, capacity plan, security plan, test environments, and any Service Operation practices.

The following table describes the type of hierarchical structures that may be used for different levels of change authorization. A degree of delegated authority may also exist within an authorization level. Formal authorization is obtained for each change from a change authority that may be a role, person or a group of people. The levels of authorization for a change should be judged by:

- Type;
- Size;
- Risk;
- Financial implications; and
- Scope.

<b>Level</b>	<b>Change Authority</b>	<b>Potential Impact/Risk</b>
1	Business Executive Board.	High cost/risk change - Executive decision
2	The IT Management (Steering) Board	Change impacts multiple services/ organizational divisions
3	Change Advisory Board (CAB) or Emergency CAB (ECAB)	Change impacts only local/ service group
4	Change Manager	Change to a specific component of an IT Service
5	Local Authorization	Standard Change

The *7 Rs of Change Management* provides a set of guiding questions that need to be answered as part of the overall assessment for changes. These questions are:

1. Who RAISED the change?
2. What is the REASON for the change?
3. What is the RETURN required from the change?
4. What are the RISKS involved in the change?
5. What RESOURCES are required to deliver the change?
6. Who is RESPONSIBLE for the build, test and implementation of the change?
7. What is the RELATIONSHIP between this change and other changes?

#### **4. Authorization or rejection of the Change.**

While the responsibility for authorization of changes lies with the Change Manager, they in turn will ensure they have the approval of three main areas.

- Financial Approval - What's it going to cost? And what's the cost of not doing it?
- Business Approval - What are the consequences to the business? And not doing it?
- Technology Approval - What are the consequences to the infrastructure? And not doing it?

When authorizing changes, it is important to consider both the implications of performing the change, as well as the impacts of NOT implementing the change. This also requires empowering the Change Manager with an appropriate level of authority, as their primary role is to protect the integrity of the IT infrastructure and the services provided to customers.

#### **5. The Change is scheduled.**

The assessment of the change will have provided an estimation of the resource requirements for delivering a successful change. Change Management will need to coordinate with Release and Deployment Management so that any activities required for the build, test or implementation of the change will be scheduled when resources are available and when they least impact on live services or business critical times.

The timing of events and eventual implementation will be communicated via the Change Schedule, and visible to the appropriate IT staff members, customers and end users. Any service disruption is documented in with the Projected Service Outage (PSO). This details any revised Service Level Agreement and service availability targets because of the events in the Change Schedule, in addition to any planned downtime from other causes such as planned maintenance and data backups.

**6. Work orders are issued for the build of the Change (but carried out by other groups).**

Change Management will coordinate with Release and Deployment to identify the groups or individuals responsible for the implementation of the various elements making up the change. This will be greatly influenced by the availability of staff and any Release policies defining how and when changes are released to the live environment.

**7. Change Management coordinates the work performed.**

Change Management plays a co-ordination role as implementation is the responsibility of others (under the direction of Release and Deployment management or from Project Management groups).

This is an oversight role to ensure that all changes that can be and are thoroughly tested. Special care needs to be taken during implementation in all cases involving changes that have not been fully tested.

Remediation Planning is a critical element during the coordination of changes. Ideally, no change should be approved without having explicitly addressed the question of what to do if it is not successful. There should be a back-out plan that will restore the organization to its initial situation, through reloading a baselined set of Configuration Items.

Only by considering what remediation options are available before instigating a change, and by establishing that the remediation is viable, can the risk of the proposed change be determined and the appropriate actions taken.

**8. The Change is reviewed.**

On completion of the change, the results should be reported for evaluation to those responsible for managing changes, and then presented as a completed change for stakeholder agreement. Major changes will require more customer and stakeholder input throughout the entire process.

The review should confirm that the change has met the defined objectives, the initiator and stakeholders are happy with the results and there have been no unexpected side-effects. Lessons learned should be embedded into future changes as part of continuous improvement. This includes whether the request should be developed as a standard change or whether another Change Model is more appropriate for future management of similar requests.

Two types of reviews are performed for normal changes:

- The review of a *service change* – immediately visible to the customer and scheduled for discussion at the next service level management review meeting; and
- An *infrastructure change* – concerned with how IT delivers rather than what IT delivers, which will be (almost) invisible to the customer.

Change Management must review new or changed services after a predefined period has elapsed. This process will involve Change Advisory Board (CAB) members, since change reviews are a standard CAB agenda item. When a change has not achieved its objectives, Change Management (or CAB) will decide what follow-up action is required.

## **9. The Change is closed.**

If the change review was satisfactory or the original change is abandoned (e.g. the needs for the change is decreased or disappears) the RFC should be formally closed in the logging system. These records will be kept for a period of time based on business, compliance, archiving or other policy requirements that have been defined.

### **5.6.8 Roles and Responsibilities within Change Management**

#### **Change Advisory Board (CAB)**

The CAB is a body that assists the Change Manager in evaluating and approving changes that have been requested. The CAB is not necessarily a physical meeting of people, but can be enabled with the use of technology such as video-conferencing. Members of the CAB are made up by IT staff, customers, suppliers and other stakeholders in such a way that all changes that impact their domain can be appropriately assessed from business, technical and support viewpoints. Typical representatives for a CAB under normal conditions are:

- The Change Manager (chairs the CAB);
- Customer representatives;
- User management;
- Application Developers/Supporters;
- Technical Experts and Consultants;
- Other Services Staff; and
- Vendors and Suppliers.



Rather than having a static list of members, the CAB should include both static and dynamic members who will attend based on the needs for the changes being discussed. Standard items for the CAB agenda include:

- Failed, unauthorized or backed-out changes;
- RFCs to be assessed by the CAB ordered by priority;
- Current and outstanding changes;
- Concluded changes for review;
- Identified process issues or improvements; and
- Advance notice of RFCs expected for review at next CAB meeting.

In order to maintain an effective meeting, CAB members should come prepared with any documentation that was produced as part of their evaluation or advice on current changes. Ultimately the final decision to authorize changes will rest with the Change Manager or whichever body or person has been delegated for this in the organization (typically the IT Director or Service Manager).

### **Emergency Change Advisory Board (ECAB)**

For Emergency Changes, a separate defined procedure and authorization mechanism should exist and be clearly documented and understood. As it will not always be possible to convene a full CAB meeting, an ECAB would be used instead; which is a sub-group that will provide the assistance for the evaluation of Emergency Changes. This may be three or four members of the normal CAB, who between them have the best overview of the potential risks and impacts of implementing the Emergency Change.

While complete testing may not be possible, it is still important to apply an appropriate level of testing as well as define any roll-back mechanisms relative to the potential risk and impact identified.

### **5.6.9 Key Performance Indicators (KPIs) of Change Management**

It is important that a balanced view of metrics is used when assessing the effectiveness and efficiency of the Change Management process. Metrics should be linked to business goals whenever practical, and to cost, availability, reliability and customer satisfaction.

These metrics include:

- Number of RFCs (Accepted/Rejected);
- Number and per cent of successful changes;
- Emergency Changes;
- Number of changes awaiting implementation;
- Number of implemented changes;
- Change backlogs and bottle-necks;
- Business impact of changes;
- Frequency of change to CIs;
- Number of disruptions (incidents and problems) caused by unsuccessful changes;
- Level of variance between estimated and actual resources and time required for changes;
- Frequency of change (by service, business area etc.);
- Number of changes (by change type);
- Number of changes recorded and tracked using automated tools;
- Staff utilization; and
- Number of unauthorized changes detected by Configuration Management.

### **5.6.10 Challenges affecting Change Management**

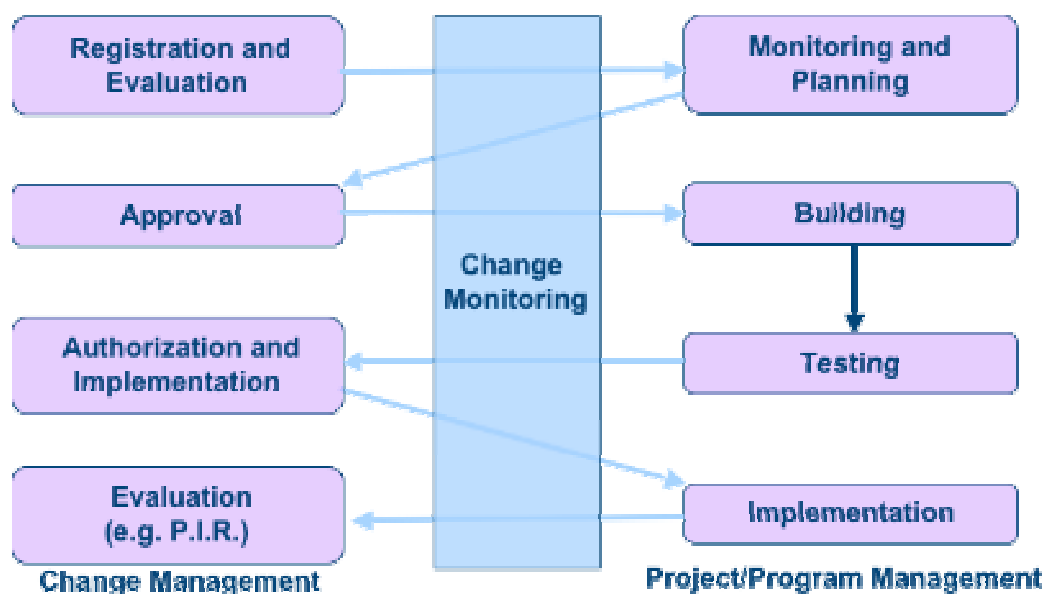
Implementing Change Management should be managed with particular sensitivity and awareness due to the often large cultural changes that will need to take place. This will need to be addressed by developing stakeholder awareness at all stages during the implementation project, including gaining feedback about their needs, expectations and ideas.

Depending on the current maturity of informal change management practices, it may be appropriate to run a pilot for Change Management, in order to identify the potential successes and failures, and to better refine and scale according to the organization's needs. Over time, the scope of Change Management should eventually include the customers and end-users just as much as IT staff members themselves; as there can be many instances where service quality can be negatively impacted as a result of changes executed by the end-user community.

Typical challenges that affect the success of Change Management include:

- Change in culture - one central process comes into place that influences everyone's activities;
- Bypassing - projects ducking the Change Management planning;
- Optimal link with Configuration Management - to execute a controlled change all data MUST be reliable;
- Commitment of the supplier(s) to the process; and
- Commitment of management.

### 5.6.11 Relationship with Project Management



#### How does Change Management work with Project Management?

- Change Management authorizes, controls, coordinates the changes involved in a given project.
- It **does not** plan, build, test or implement.
- Change Management is also concerned with Remediation Planning to ensure that each RFC has a fallback / rollback plan.

**5.6.12 Typical Contents of Change Documentation**

<b>Attribute on the change record</b>	<b>RFC</b>	<b>Change proposal (if appropriate)</b>	<b>Related assets/CIs</b>
Unique Number	✓		
Trigger (e.g. to purchase order, problem report number, error records, business need, legislation)	✓		
Description	Summary	Full description	
Identity of item(s) to be changed – description of the desired change	Summary	Full description	Service (for enhancement) or CI with errors (corrective changes)
Reason for change, e.g. Business Case	Summary	Full justification	
Effect of not implementing the change (business, technical, financial etc.)	✓		
Configuration items and baseline versions to be changed	✓	Affected baseline / release	Details of CIs in baseline / release

Contact and details of person proposing the change	✓		
Date and time that the change was proposed	✓		
Change category, e.g. minor, significant, major	Proposed		
Predicted timeframe, resources, costs and quality of service	Summary / reference	Full	
Change priority	Proposed		
Risk assessment and risk management plan	Summary / reference	Full	
Back-out or remediation plan	Possibly	Full	
Impact assessment and evaluation – resources and capacity, cost, benefits	Provisional	Initial Impact	✓

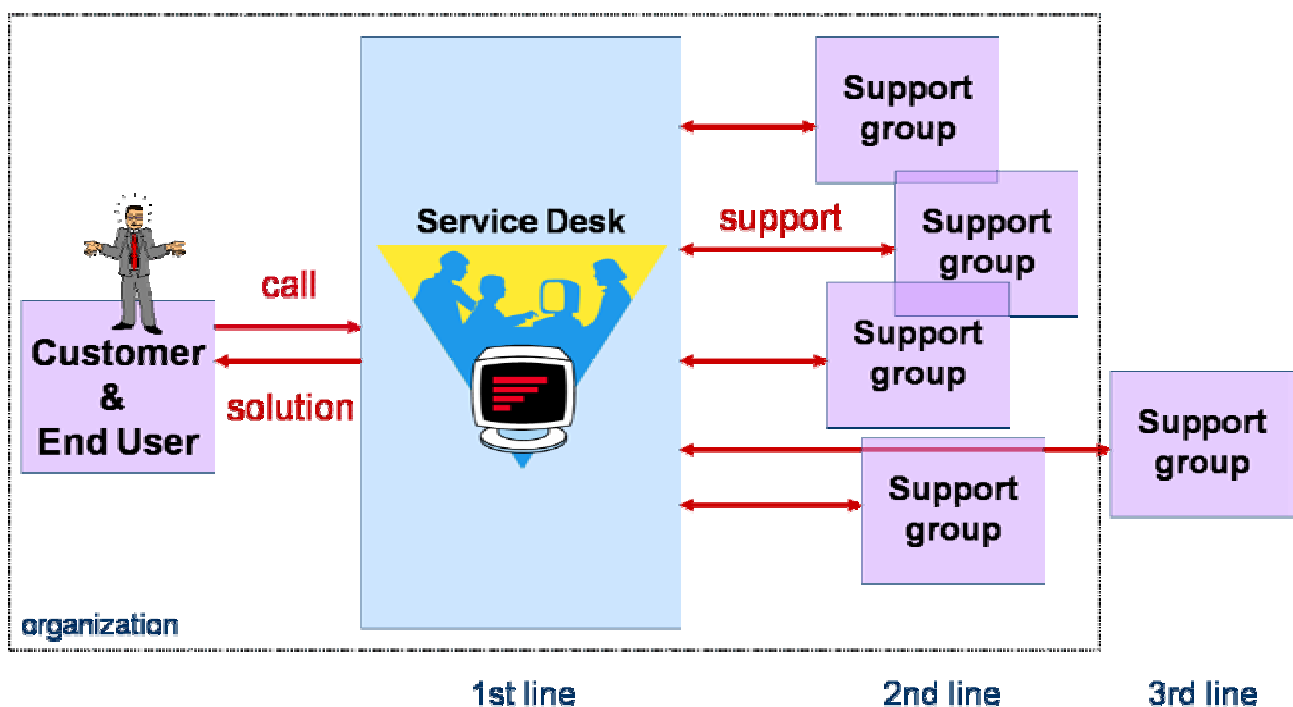
Would the change require consequential amendment of IT Service Continuity Management (ITSCM) plan, capacity plan, security plan, and test plan?	✓		Plans affected
Change decision body	✓		
Decision and recommendations accompanying the decision	✓		
Authorization signature (could be electronic)	✓		
Authorization date and time	✓		
Target baseline or release to incorporate change into	✓		
Target change plan(s) for change to be incorporated into	✓		

Scheduled implementation time (change window, release window or date and time)	✓		
Location / reference to release / implementation plan	✓		
Details of change implementer	✓		
Change implementation details (success / fail / remediation)	✓		✓
Actual implementation date and time	✓		
Review date(s)	✓		
Review results (including cross-reference to new RFC where necessary)	Summary		
Closure	Summary		

## 5.7 Service Desk

**Objective:** To support the agreed IT service provision by ensuring the accessibility and availability of the IT organization and by performing various supporting activities.

A Service Desk is a functional unit that acts as the primary point (first line) of contact for the end-user community for all incidents, requests and general communications that arise. It plays an essential and valuable role for any organization, contributing significantly to the satisfaction of users and the overall impression of the IT organization. Depending on the type of business, services and technology supported, the exact size and physical organization of the Service Desk will vary from a small centralized team to a diverse range of teams in multiple locations and time zones.



### 5.7.1 Goal and objectives

The primary goal of the Service Desk is to support the agreed IT service provision by ensuring the accessibility and availability of the IT-organization and by performing various supporting activities. Other objectives include:

- To act as a single point of contact for all user incidents, requests and general communication;
- To restore 'normal service operation' as quickly as possible in the case of disruption;



- To improve user awareness of IT issues and to promote appropriate use of IT services and resources; and
- To assist other the other IT functions by managing user communication and escalating incidents and requests using defined procedures.

### **5.7.2 Benefits**

While many organizations have already seen the justification for the creation of a Service Desk team(s), in many cases the business case for the improvement fail to gain support from various levels of management. As discussed earlier, the needs and requirements will vary significantly for each organization; however, the typical benefits gained through the implementation/improvement of a Service Desk function include:

- Improved customer service perception, and satisfaction;
- Increased accessibility through the use of a single point of contact;
- Better quality and speedier turnaround of requests;
- Improved teamwork and communication;
- Better managed infrastructure and control; and
- Improved usage of IT resources.

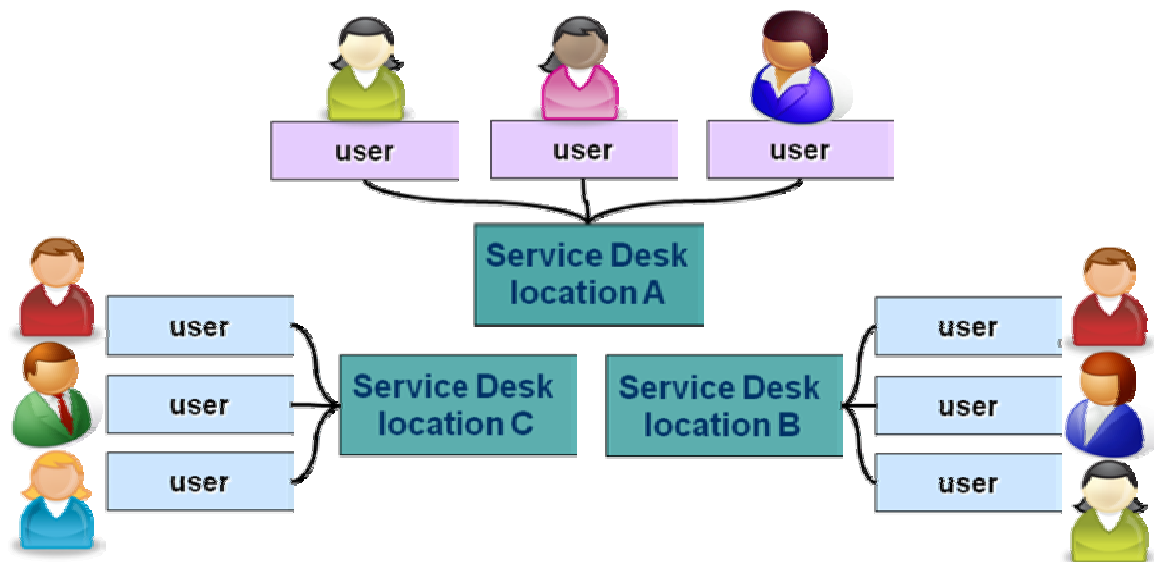
### **5.7.3 Service Desk organizational structures**

Many factors will influence the way in which a Service Desk function will be physically structured, such as the location, languages and cultures of end users, diversity in services and technology supported and the objectives governing the implementation of the Service Desk such as improved satisfaction or reduced operating costs.

The following is some of the main options chosen when implementing a Service Desk function:

## Local Service Desk

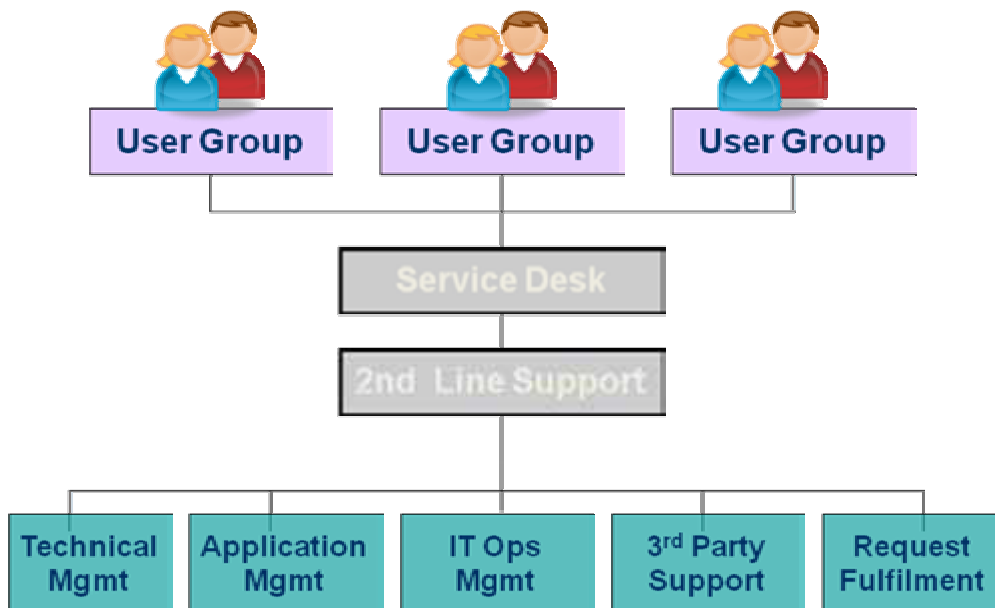
A local Service Desk structure is where the Service Desk is co-located within or physically close to the user community it serves. This may aid in communication and give the Service Desk a visible presence which some users may like. It may, however, be inefficient and expensive as a result of having multiple Service Desks operating.



Benefits of a Local Service Desk structure	Disadvantages of a Local Service Desk structure
<ul style="list-style-type: none"> <li>• Local and specific-user knowledge</li> <li>• Ability to effectively communicate with multiple languages</li> <li>• Appropriate cultural knowledge</li> <li>• Visible (and physical) presence of the Service Desk</li> </ul>	<ul style="list-style-type: none"> <li>• Higher costs for replicated infrastructure and more staff involved</li> <li>• Less knowledge transfer, each Service Desk may spend time rediscovering knowledge</li> <li>• Inconsistency in service levels and reporting</li> <li>• Service Desks may be focused on local issues</li> </ul>

## Centralized Service Desk

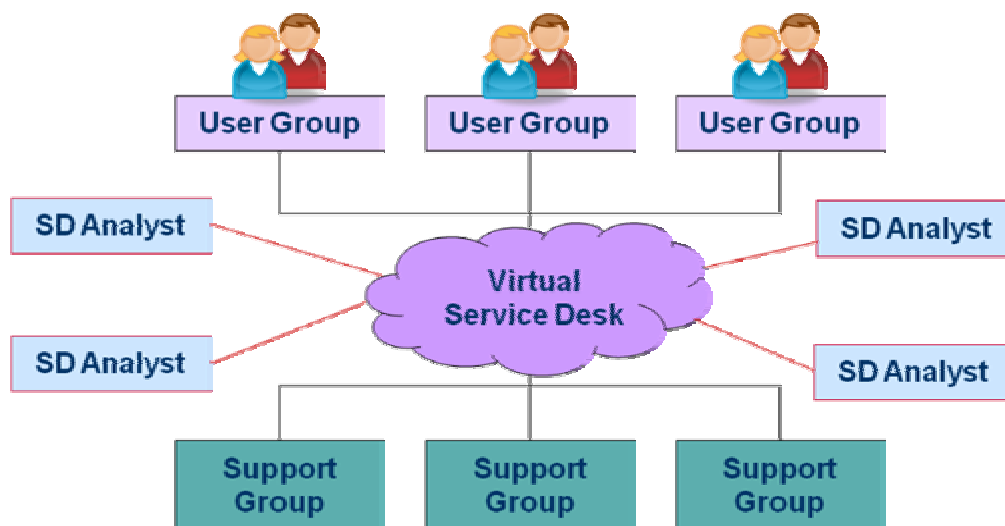
A centralized structure uses a Service Desk in a single location (or smaller number of locations), although some local presence may remain to handle physical support requirements such as deploying, moving and disposing of user workstations. This could be more efficient, enabling less staff to manage a higher volume of calls, with greater visibility of repeat incidents and requests.



Benefits of a centralized Service Desk structure	Disadvantages of a centralized Service Desk structure
<ul style="list-style-type: none"> <li>• Reduced operational costs</li> <li>• Improved usage of available resources</li> <li>• Consistency of call handling</li> <li>• Improved ability for knowledge sharing</li> <li>• Simplicity for users (call one number) to contact the Service Desk</li> </ul>	<ul style="list-style-type: none"> <li>• Potentially higher costs and challenges in handling 24x7 environment or different time zone</li> <li>• Lack of local knowledge</li> <li>• Possible gaps in language and culture</li> <li>• Higher risk (single point of failure), in case of power loss or other physical threat.</li> </ul>

## Virtual Service Desk

A Virtual Service Desk through the use of technology, particularly the Internet and the use of corporate support tools, can give users the impression of a single, centralized Service Desk when in fact the personnel may be spread or located in any number or type of geographical or structural locations.



Benefits of a virtual Service Desk structure	Disadvantages of a virtual Service Desk structure
<ul style="list-style-type: none"><li>• Support for global organizations</li><li>• 24x7 support in multiple time zones</li><li>• Reduced operational costs</li><li>• Improved usage of available resources</li><li>• Effective matching of appropriate staff for different types of calls</li></ul>	<ul style="list-style-type: none"><li>• Initial cost of implementation, requiring diverse and effective voice technology</li><li>• Lack in the consistency of service and reporting</li><li>• Less effective for monitoring actions of staff</li><li>• Staff may feel disconnected from other Service Desk staff</li></ul>

## ***Follow the Sun***

Some global or international organizations will combine two or more of their geographically dispersed Service Desks to provide 24-hour follow-the-sun service.



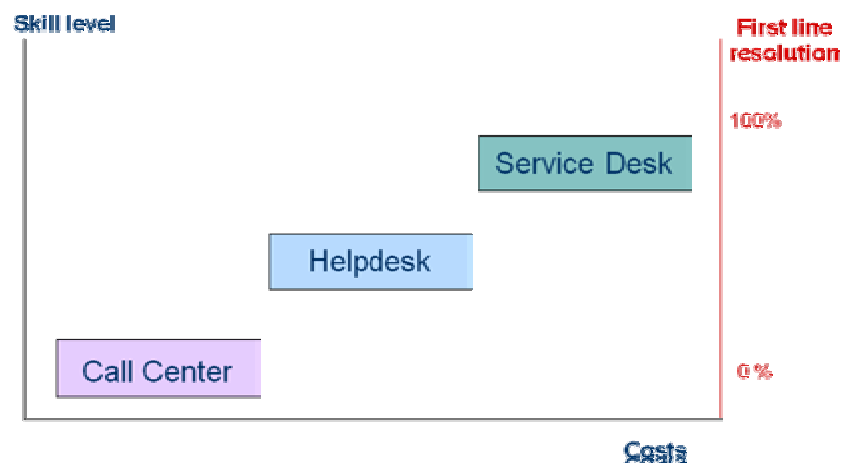
<b>Benefits of a 'Follow the Sun' Service Desk structure</b>	<b>Disadvantages of a 'Follow the Sun' Service Desk structure</b>
<ul style="list-style-type: none"><li>• Support for global organizations</li><li>• 24x7 support in multiple time zones</li><li>• Improved quality of service</li><li>• Improved customer/user satisfaction</li><li>• Effective knowledge sharing and high level visibility of distributed infrastructure</li></ul>	<ul style="list-style-type: none"><li>• Typically higher operating costs</li><li>• Cost of required technology</li><li>• Challenges in using single language for multiple regions when recording knowledge, workarounds, Known Errors etc</li></ul>

### 5.7.4 Service Desk Types (skill levels)

Depending on the requirements defined for the Service Desk, organizations will need to consider what skill level is appropriate for the Service Desk and the support it will offer. This skill level can be defined in many ways, but most often is associated with the first time resolution achieved for calls, incidents and requests made to the Service Desk by users.

The 3 types of Service Desks are:

- **Call Centre:** Responsible for *handling/logging* of large volumes of calls;
- **Help Desk:** Responsible for **managing** and *co-ordinate* incidents; and
- **Service Desk:** Responsible for **managing** incidents and requests, and also provides a wide variety supporting services (e.g. supporting Human Resources).



### 5.7.5 Service Desk staffing

One of the most challenging issues facing a Service Desk Manager is that of staffing, with many organizations finding it increasingly difficult to acquire and retain quality employees. Additionally, determining appropriate staff levels can be difficult, with call rates being very volatile and dynamic. The following section will describe some of the issues involved with staffing a Service Desk.

#### **Hiring Service Desk staff**

Most Service Desk Managers will have a list of key competencies or selection criteria that is used when hiring new staff members. Depending on the type of Service Desk that has been implemented and what types of technologies are being supported, these criteria will vary; however, typical skills required include:

- Communication skills;

- Technical knowledge;
- Business understanding;
- Diagnosis and analysis skills;
- Understanding of the role/value of processes and procedures; and
- Typing skills.

Communication skills are ultimately the most important as they will need to be able to deal effectively with a wide-range of people and stressful situations.

### ***Service Desk staffing levels***

The number of staff employed on the Service Desk is dependent on the needs of the business, the objectives/goals defined and a range of other important criteria including:

- Business budget;
- Customer service expectations;
- Size, maturity, design, complexity of the IT Infrastructure & service catalogue;
- The number of customers and users to support;
- The volume of requests, incidents and general communication required;
- Period of support cover required;
- Workload pattern of requests;
- SLA definitions in place;
- Level of training required; and
- Support technologies available.

### ***Super Users***

‘Super Users’ are often useful positions to be appointed across an organization to act as liaison points with the IT organization and the Service Desk in particular. Super Users can be used in a number of ways such as:

- To assist in general communication between the Service Desk and users;
- To filter requests and issues raised by the user community; and
- To assist in user training.

While Super Users can be a valuable resource if they are properly coordinated, the following rules should be in place:

- Roles and responsibilities are clearly defined;
- Escalation channels are defined;
- Standard support processes defined and used; and
- All requests recorded and maintained consistently.

### ***Staff retention***

To ensure a balanced mix of experienced and newer staff, Service Desk Managers should use a number of methods and incentives to retain quality staff and to avoid disruption and inconsistency in the quality of support offered.

Some ways in which this can be done include:

- Recognition of staff achievements contributing to service quality;
- Rotation of staff onto other activities (projects, second-line support etc.);
- Team building exercises and celebrations; and
- Promote the Service Desk as a potential stepping stone for staff to move into other more technical or supervisory roles (after defined time periods and skills achieved).

### **5.7.6 Key Performance Indicators (KPIs) for the Service Desk**

To evaluate the true performance of the Service Desk, a balanced range of metrics should be established and reviewed at regular intervals. Especially dangerous is the tendency to focus on “average call time” or “number of calls answered” metrics which can mask underlying issues with the quality of support provided.

Some of the typical metrics reviewed when monitoring the performance of the Service Desk include:

- The number of calls to Service Desk (broken down by type, time of day and day of the week);
- First-line resolution rate;
- Average Service Desk cost of handling any incident or request;
- Number of knowledgebase articles created;
- Number or percentage of SLA breaches;
- Call resolution time;
- Customer satisfaction (surveys); and
- Use of self help (where exists).

### **5.7.7 Outsourcing the Service Desk**

Although fairly common, there are potential risks that can be introduced when outsourcing an organization's Service Desk. When reviewing the potential for this to occur, Service Managers should consider the following items when developing contracts to reduce these risks:

- Use of your own Service Management tool, not theirs;
- Retaining ownership of data;
- Ability to maintain required staffing levels;
- Agreements on reporting and monitoring needs;
- Proven up-to-date procedures;
- Agreed and understood support needs; and
- Engage contract specialists for assistance.



## ***6 Virtualization and Storage Management***

### ***6.1 The Heart of Cloud Computing: Virtualization***

The success of cloud computing would not be possible without the advances in virtualization. The basic concept is to have one resource pretend to be another. Its implementation is applicable to many areas of computing. Virtualization allows a computer to become several different computers, or one software package to emulate another.

This concept is critical to cloud computing and storage as a service. In a storage-as-a-service solution, the user is provided a logical space for data storage. In a virtualization system, this logical space is mapped to an actual physical location in a concept called address space remapping, the actual form is dependent on the granularity of the mapping. In a block-based storage environment, the method is to address each block of information with a logical unit identifier (LUN). The address space mapping is between the logical disk or virtual disk (vdisk) and a logical unit presented by storage controllers.

To handle this address space mapping, the virtualization software is responsible to provide a consistent view of all the mapping information attributed to the storage. This view is stored as a mapping table and called meta-data. This meta-data is used to re-direct input/output (I/O) requests. The virtualization software will receive a request for information located on the vdisk and will route the request to the physical disk location.

The primary driver for virtualization is savings. By virtualizing under-utilized systems onto a single server, businesses can realize savings because less power and space is being consumed, as well as reduced efforts in cooling and administration. Virtualization allows an operation system and applications to be migrated to a new server as a way to load balance. This activity is called live migration. In a development environment, virtualization allows greater ability to troubleshoot problems. Since multiple operating systems can virtually run on a single system, when one system crashes due to a bug, the entire system does not crash allowing the developer to debug the problem.

## **6.2 Defining Virtualization**

SNIA defines virtualization in two ways:

- The act of abstracting, hiding, or isolating the internal functions of a storage (sub)system or service from applications, host computers, or general network resources, for the purpose of enabling application and network-independent management of storage or data; and
- The application of virtualization to storage services for the purpose of aggregating functions or devices, hiding complexity, or adding capabilities to lower level storage resources.

In short, the definition of virtualization is “abstraction in detail.”

A typical model of a storage solution has four layers. The first layer is the physical storage devices themselves. Above this is the block aggregation layer where the logical definitions of the devices as well as servers, network components and clients fall. The third layer is handles the database and file systems. And the application layer is the processes and controls used to handle the data. Virtualization allows each layer to create the previous form into a new one. For instance, several storage devices can appear as one in the block aggregation layer.

## **6.3 Why Virtualize?**

The primary reason for virtualizing the storage network is to resolve problems that have been inherent in managing storage area networks. Sooner or later, companies will have to deal with the same basic problem: multiple systems being used at every level of the network stack. The diversity is caused over time: as the business grows, new systems are added. As systems degrade, components are replaced. As technologies emerge, they are adopted into the infrastructure. In the end, the storage infrastructure becomes a hodge-podge collection of complex systems and controls. The result is a growing risk to achieving business goals. Virtualization can improve several areas of managing a SAN, including:

- Providing cost-effective ways to eliminate single points of failure in the SAN;
- Improving quality of service by managing and improving performance in real time;
- Providing cost-effective solutions to disaster recovery and data archiving;
- Improving utilization across the SAN; and
- Improving scalability and flexibility on the network.

Essentially, virtualization reduces complexity. In doing so, a virtualized network is easier to manage. Easier management means less cost to administering a complex solution. Enhanced scalability, flexibility, and utilization provides further cost savings. Automation is an additional functionality available through virtualization. Without it, resolving problems would require an understanding of the specific characteristics of each component in the solution.

## **6.4 What Can Be Virtualized?**

The SNIA taxonomy lists five different types of storage virtualization: disk, tape (media and drive), file system, file, and block.

- Disk virtualization is one of the oldest forms of storage virtualization. The physical form of a magnetic disk is a compilation of cylinders, heads, and sectors. Each disk is different based on the numbers of cylinders, heads, and sectors; which changes the capacity of the disk. In order to read or write using the disk, some form of addressing is required. To complete this, addressing the exact physical property of every magnetic disk would have to be known – an impossibility. Virtualization is completed by the disks' firmware to transform the addresses into logical blocks for use by applications and operating systems, called logical block addressing (LBA). As magnetic disks are used, some blocks may not be able to store and retrieve data reliably. Disk virtualization allows the magnetic disk to appear defect-free, releasing the operating system to focus on other processes.
- When data is read or written to physical tape media, the transmission can fluctuate due to the network or busy clients on the network. The ideal situation is to have a single stream of data. When fluctuations occur, the tape drive reverts to a start/stop mode where a small amount of data is written; the tape is stopped and rewind before another small amount of data is written. This mode increases recording time and wear on the tape media. Variances like this result in only 15-30 percent of tape media capacity to be utilized. Tape media visualization uses online disk storage as a cache to emulate reading and writing. This disk acts a buffer to eliminate the fluctuations and variances present when writing directly to the tape media. Tape media virtualization also emulates large numbers of small volumes while keeping the data on disk. The data is then writing to tape in a streaming pattern.
- Tape drives used to be dedicated to single servers and then connected to multiple servers. Now tape drives are connected to a server and the activities involving these devices are performed across the network. Unfortunately, a potential risk is two servers writing to the same tape drive at the same time. Virtualized tape drives can make a single drive appear as several virtual drives which are assigned to individual servers as dedicated resources. When a request is made, the virtualization software (tape broker) reserves and maps a physical tape drive to the virtual drive and completed the operations. The physical drive is then placed back

into the environment to be used by another virtual drive. Another virtualization form is the use of RAIL (Redundant Array of Independent Libraries). This technique allows several tape libraries to be emulated and distributed as a single virtual library for use by applications.

- Files systems virtualization provides file systems to multiple clients regardless of the operating systems on those clients. Using a networked remote file system such as NFS or CIFS, this abstraction demonstrates the most important feature of storage virtualization: location transparency. Another form of file system virtualization assists in database management. Most data for databases are located on raw disk drives to maximize performance but are cumbersome to manage. Putting this data into a file system makes the data easier to manage but causes problems in performance. Database virtualization combines the best of both worlds.
- Hierarchical Storage Management (HSM) identifies rarely used data and migrates it to inexpensive secondary storage. From the client perspective, the file still exists and there is no need to know the physical location of the file. This is a form of file virtualization.

Block virtualization, or block aggregation, creates a single logical device out of several physical devices. The intent is to have applications see a new virtual disk with a larger range of block addresses; essentially tricking the application into thinking it has more memory than it truly has. The storage customer is concerned with three things: capacity, performance, and availability. When the customer demands more, functions of block virtualization can fulfill these requirements easily and cost-effectively.

## ***6.5 Where Does Virtualization Happen?***

Every system has a level of block virtualization implemented into the architecture. Recent changes in managing networks have made block virtualization a key technology. In fact, block virtualization has become so prominent in recent implementations or new services and transformations of old systems that it can be found at every layer of the storage solution.

To understand the possibilities of virtualization, it is important to understand the simplest storage operation: an I/O request from application to storage. An application makes a request to the operating system to read or write to a file. The I/O request is transformed into a logical block address by going through the file system. The logical address is converted into a physical disk address. This conversion can happen at the host layer, the network layer, the storage layer, or a combination of the three. When the conversion to physical address is complete, the disk is not accessed and the results move back through the steps to the application. Each layer of the storage infrastructure now becomes a stack of technologies responsible for the conversion of the I/O request. Virtualization can

happen at any one of these stacks. Most virtualization approaches happen at the host, network, and storage stacks.

Host-based virtualization is typically attributed to logical volume managers (LVM) that can be found on any computer from a desktop to a data center. LVMs are a standard part of the operation system with more sophisticated LVMs available from third-party vendors. Since file systems are also found on the host, more efficient capacity management is possible by tightly converging the two. The downside is that LVMs are server-centric requiring storage provisioning to be performed on each host. In a large-complex environment, managing LVMs on each server can be labor intensive, though some third party vendors have developed cluster volume managers.

Storage-based virtualization does not require a specific type of host. Therefore, the storage array can support various operating systems and applications. Some of the more common virtualization techniques found on the storage array are RAIDs (Redundant Array of Independent Disks), snapshots, LUN (Logical Unit Number) masking, and mapping. In many solutions, the virtualization at the host and storage stacks are combined to gain the flexibility of host LVMs and the performance of storage RAIDs.

Network-based virtualization allows a truly flexible SAN solution to be realized. It can combine several LUNs from one or more arrays into a single LUN or break a single LUN from an array into smaller virtual LUNs. Synchronous and asynchronous replications are both possible through virtualization. Other functions at the network stack include caching, storage on demand and quality of service (QoS) functions.

## **6.6 How Does Virtualization Happen?**

Integration into the storage infrastructure happens in two ways: in-band virtualization and out-of-band virtualization. In-band virtualization speaks to the addition of a SAN appliance to the network, specifically between the host and storage. A SAN appliance allows information and data to pass through it. To the host, the appliance looks like storage array. To storage, the appliance looks like a host. The purpose of the appliance is to control access and essentially routing I/O requests through the network or acting as a storage firewall. The use of an appliance removes any requirements for agents on the host, simplifying implementation and management of the SAN.

Out-of-band virtualization has the appliance sitting outside the data path of the host and storage. The server must contact the appliance before accessing the storage; therefore the host must know that the storage is virtualized. This lack of transparency forces additional agents and controls to be added to the host server, creating potential performance issues. Implementation into an out-of-band solution is far more difficult.

Several functions are fulfilled by SAN appliances: managing the storage pool, volume configuration, controls, identifying devices, and I/O error handling.

## ***6.7 On the Road to Storage Virtualization***

Though different companies will have different virtualization implementations, there are few general commonalities that each company should look at to determine how they get from a conventional storage environment to a virtualized environment. In any situation, cost effectiveness and performance are major considerations for any step in the implementation process.

The first common concern for a company moving to a virtualized environment is to determine if any existing storage equipment can be used. Do they have an existing SAN? Can that equipment or any other work in a virtualized SAN? Can firmware upgrades resolve any compatibility issues? What architectural considerations need to be applied to use the existing equipment?

For any major migration a full and complete backup must be conducted, though recovery is very unlikely. Data can be impacted by the method of virtualization. Out-of-band virtualization will not touch the data at all. In-band virtualization requires a solution that will layer the virtualization process onto the existing LUN without touching the partition table or file systems. And even if it the solution can do this, some set of the advanced virtualization functions will not be possible. An alternative is to use host-based mirroring to adapt the existing LUN into a fully virtualized LUN.

The actual implementation is highly dependent on the methods used, specifically out-of-band and in-band approaches. In the out-of-band approach, each application server must have software installed to allow it to be aware of the virtualization layer. In the in-band approach, the SAN appliances need to be added to the environment and recognized as storage by the hosts and as hosts by the storage units. This is a simplified explanation of implementing the virtualization engine. Once all application servers are connected to the SAN and the engine is in place, the next step is to ensure the data is migrated from the direct attached storage to the SAN storage, then removing the direct attached storage. An additional step of creating a system redundancy is highly recommended.

All attempts to virtualize an IT environment should be done by an expert. There are many considerations in place at each step of the implementation process that require extensive knowledge about the capabilities of virtualization.

## ***6.8 Improving Availability Using Virtualization***

When the initial implementation to a virtualized environment is complete, the work is not yet complete. Now improvements to the system should be identified and implemented. One concentration is on improving availability. The goal here is to reach a very high level of availability - that is 99.99% uptime. In order to do this, the system should always have a redundant instance of major control factors, such as switches, appliances, and virtualization engines. This effort in redundancy will ensure that if a major component fails, an alternative route exists to continue operations.

To obtain this redundancy a dual-engine or clustered-engine approach is used with a great degree of transparency to the host. Several techniques are found in topologies geared towards high availability. The active-passive technique has two virtualization engines in place, the first is active at all times, while the second is used only in the event of engine failure or overload. The active-active technique has the same two engines sharing the normal workload. When one engine fails, the second simply takes the additional workload until the first is up and running again. Within a clustered approach instead of two hosts, multiple hosts are utilized. This is called an N-way design: when one host fails, the system of  $N - 1$  hosts still remain.

No matter the approach used, a complete virtualization solution geared towards high availability will have redundancies at every layer of the design: multiple host support, multipathing, replication capability, cluster file system capability, multiple virtualization engines and subsystems. In a virtualized environment, there is little worry if a single system or server fails, as there is another in place to compensate allowing the solution to be available at all times.

## ***6.9 Improving Performance through Virtualization***

Even though most storage virtualizations will enhance the storage performance of a SAN, it is prudent for the IT professional to consider better ways to improve performance. There are so many virtualization approaches that planning them out will provide certain advantages and disadvantages to cost and performance that need to be balanced out. For instance, focusing on network-based virtualization relieves the burden of the task from the host and storage stacks. Out-of-band virtualization forces the host to carry the burden as it must be present to the virtualized environment. Most performance enhancements of the virtualization system can be done on any stack of the solution.

The two most prominent performance enhancements are striping and caching. Striping is

a method of writing data across storage disks simultaneously. Though a virtualized system, the striping can be done across disk within an array or multiple arrays. The more physical devices the data is striped, the better the throughput and I/O per second. Caching allows data to be stored in logical memory before being written in physical memory. It is particularly advantageous in environments where different components have a large difference in processing speed.

Most computer systems and storage arrays already have caching functions available to them. Network-based virtualization provides additional caching principles such as moving frequently accessed data closer to the host layer. Mirroring and dual pathing are two additional network-based virtualization techniques that allow a balance to the workload. Though these techniques may have originated because of improvements to availability, they also provide performance enhancements.

## ***6.10 Improving Capacity through Virtualization***

Of all the disciplines within storage management, capacity management has the most challenges associated to it. The reason for this is because when all is said and done, adding physical storage to the system is still a highly manual, highly complex activity with much room for error. The concerns of capacity occur almost immediately and can get progressively worse as time goes by. Capacity managers are constantly dealing with storage disks that are too small or too large to support the intended application, capacity is not available where it is needed the most, and over time, disks will fill up and run out of storage space. Trying to add new storage resources take time, mostly in planning, simply to ensure that nothing is broken; not to mention that a host will have different requirements for different volumes of storage.

Before virtualization, a common practice for resolving capacity issues was to attach larger storage capacities than required to the host environment. This would reduce the number of times capacity was added, reducing the risk in error but decreasing the utilization of the capacity in the environment.

Storage virtualization provides the opportunity to pool storage, simplify allocation, expand LUNs, and automate capacity controls. Pooling storage simply allows the host to see one single storage entity when in fact the physical space belongs to several storage devices. In a traditional storage solution, allocation management is done at each host or through different segments of the network. In a virtualized environment, the entire system is controlled by a common management interface allowing allocation and zoning to be simplified. LUNs are restricted by the physical limits of the disk drive, yet virtualizing LUNs allow those limits to be removed so that it appears that more LUNs are available than there really are. Virtualized SANs allow a greater amount of automation to capacity



management functions.

Expanding LUN is starting to become a common solution to adding capacity to the environment. The conventional approach has the system administrator, or agent software, detect where free space has fallen below a specified threshold. Additional capacity is assigned appropriately using a new LUN or resizing the existing LUN. Another approach uses a “sparse volume,” a virtual device that the host can see. However, physical storage is not used until data is written to the virtual LUN.

### **6.11 Business Value for Virtualization**

Organizations are dealing with rapid data growth and expansion in the range of applications. These conditions have a significant impact on how they setup their datacenters. Virtualization allows an increase in IT asset utilization, availability, and reduction in administration management. The improper deployment of storage assets to virtualized servers can result in unplanned impacts to the storage environment, including long and expensive projects to re-architect the network and migrate storage, and problems with provisioning. Therefore, for any company looking to implement some virtualization to their servers and the network the servers are attached to should consider virtualization of the storage area network. By doing so, the company benefits by a tiered storage solution that is fast, flexible and efficient, data migration and movement of applications that are disruptive, efficient snapshot and replication of images and data for backup, test, and configuration, and the availability and rapid recovery of applications and data.

The emergence of virtual machines in the early 2000s has become one of the fastest growing technologies in IT. Lack of capital resources, need for additional capacity, and difficult power and cooling issues have been significant reason for virtualizing servers and reclaim excess capacity. Production-level consolidation has been a major reason lately for virtualization.

The biggest benefits for virtualization are:

- Standardized configuration of storage devices. A virtual environment requires a minimum level of standardization on the storage components. With a consistent configuration throughout the drives, it is easier to drive availability uptime in the environment. The virtualization software allows all operating systems to map to the same file system.
- Easier migration. Data migration can result in lengthy downtime. Virtualization can reduce this downtime.

## ***7 Applying Cloud Computing to Project Management***

### ***7.1 Benefiting from Cloud Computing***

The benefits of cloud computing include quicker implementations and lower costs while providing greater scalability, adaptability, and reliability to the environment. Especially in terms of web applications, a solution can be available at any time, in any location on the planet, by any person. As use of the application increases or decreases, the cloud solution adjusts accordingly.

The previous pages showed the benefits of project management applied to cloud computing. The following pages will explore the benefits of cloud computing to project management. Some of the ideas explored here have solutions found within the cloud environment, especially in Independent Software Environments (IDEs) used for software development.

Cloud computing solutions exploit a number of distinct technologies:

- Virtualization is a technique used to create an abstract rendering of a physical component. For instance, storage volumes can be virtualized. With a virtual volume, the storage capacity can be increased, partitioned without partitioning the physical drive, or merging the volume with other storage devices to create a single large volume. Virtualization allows for increased scalability, reliability, and portability to occur.
- Web 2.0 is a perspective change for how to use the web. Often referred to as a technology, Web 2.0 provides a number of options that make web applications just as rich as desktop applications.
- Open source is a concept supporting total collaboration between developer and user. Often attributed to application development, the concept is also applicable to any number of solutions where user knowledge is utilized extensively to generate the product that the user is concurrently using. The idea ultimately encourages development in production.

Seven benefits are available to project management when adapting cloud computing capabilities:

- The ability to interlink distinct components together as a complete whole. For instance, linking activity with cost, time, risks, quality, and resources. Many project management solutions already do this to some extent; however, with cloud computing the options can be expanded to include historical information, feedback, and particular notes;

- To support communication specifically with stakeholders requiring performance reports, the use of subscription services allow distribution of information to have greater potential and customizable by the stakeholders themselves;
- Many aspects of project management can be present at any given time and with an application, the possibility of an interface that allows each aspect to deliver is highly probable with web applications;
- One of the risks to project management is ensuring that all the right players are involved in the project to ensure that everything that needs to be done on the project has been planned. With cloud computing, participation by all parties is possible;
- The core tenet of cloud computing is the ability to focus on components as parts of the whole allowing an individual component to be added, removed, expanded, and reused without impact the whole;
- The capabilities of the web allow portability. First of all, a web application can be accessed anywhere, at any time, by anyone with just a browser and a connection to the Internet. Second, with the business logic of the application on an Internet server, nothing is lost between multiple users; and
- A main driver of web applications and cloud computing is an improved, richer user experience. Improving the experience of individuals on a project will encourage communication, productivity, and morale.

## **7.2 The Ease of Linking (Hyperlinks)**

Looking at the construction of the web, one will see a series of documents in place, basically organized through the URL hierarchy, and with some programming scripts to provide some interest. Though all of these components make up an effective website and web application, the one construct that makes the web, the web has always been the ability to hyperlink. Hyperlinks are connections from an original page to a source document. By clicking on a hyperlink, a user can be moved from one page to another with just one click. By hovering over the link, information can be made available to the user. Through searches, multiple hyperlinks can be generated on a single page based on relevancy for the users to peruse.

Hyperlinks can be used to transform basic pieces of data into rich volumes of information by creating relationships. In the back end of an application a database is usually found that holds data and created relationships with that basic data, such as a person's name and their job position. Using the web and hyperlinks, building relationships can be expanded to incorporate multiple databases. So in addition to a name and a job position, the person's personal profile and the job description for the position can be connected. This connection can be presented as links in a single page or the full or partial renderings of the sources.

There are several perspectives that are prevalent in project management, each with some level of importance to the users. But the perspectives all rely on the need to understand the relationships with different pieces of data. Most project management solutions handle the basic relationships, such as the relationship between activity, cost, schedule, and resources. Some solutions even provide notes for referencing. However with hyper linking, any document or parts of a document can be referenced allowing risks and risk responses. Contractual terms, historical information, procurement proposals and supplier information, quality requirements, and the like can be linked to a single activity. By having the activity show all related information pertaining to it, all participants on the project can better understand and manage the work expected. The greatest aspect of this solution is that the information can exist in multiple databases and separated locations.

Even connections with existing and past projects can be made. In situations where similar projects are performed, like building a house with similar designs, connections can be made very easily. Why make these connections? Each constructed building can have its unique attributes: location, features, options, even contractors can be different. Using the designs as the starting point, a construction company can identify and compare the differences between buildings using the same design. The opportunities available are to assist in understanding requirements, risks, and unplanned threats to a particular project. For instance by referencing contractors familiar with the design, a replacement contractor can be found quickly if the original supplier is unable to do the job.

What is linked and to what extent the linkages are made is left only to the imagination of the performing organization.

### ***7.3 Subscribing to Success (Blogging)***

Imagine compiling information into the next status report, categorizing it and posting it out to the web. Without any extra effort, a notification of the report is sent to every stakeholder interested in the project. Within minutes, responses to the report start to show up. The next day, all the responses are compiled for the project team meeting to review and answer. This example is just one of the ways blogs can assist in project management. Blogs utilize a subscription web service to ensure that people who want to be informed, are. Used by a number of columnists, both professional and amateur, a blog is a mechanism used to post information; with logic to identify who is interested in the information for a notice to be sent. The identity of interested parties was provided by the people themselves when they “subscribed” to the blog. Since a blog is basically a web page, the format of the report can include anything that a typical web page can include.

Imagine again that the project is going along nicely. Coming into work early, the computer is turned on and waiting for you is a list of all the project activities that are currently being performed, the ones that were completed the previous day, and the ones that are late are highlighted. As a project manager, the list includes all the activities of the project. As a

project team member, the list includes only those which are assigned to them or have a dependency on. The mechanism used to compile this information from a project plan is the same used by blogs.

At the heart of the blogging feature is the RSS technology. Really Simple Syndication (RSS) allow a person to subscribe to a web page so that when it changes, they are notified. The technology is not attached to the webpage though, but to the link between the user and the web page. A RSS feed allows these links to exist. Aside from blogging, the RSS technology has made another significant contribution: permalinks, or permanent links. With the permalink feature, users can link to other user's websites, link to individual comments on a page, use "trackbacks" to other links to the page by other others, or respond to those other links, thus creating a two-way link. In project management, permalinks can provided a variety of associations between items. For instance, dependencies within a project can use permalinks to make two way associations possible, so that the project team can identify the activities and their dependencies, both dependent upon as well as depended for.

Hyperlinks provided a way to show one way relationships between two items. RSS has allowed two way associations to be made. The use of RSS in blogging solutions has provided a new platform for gaining information within business. For project management, it's a simple way of keeping stakeholders informed of what is going on with the project.

### *Interfacing with the Project (Mashups)*

During project management, the number of items to be monitored and controlled is considerable. For the project manager, the concerns to be watchful of include:

- Is the scope of the project intact? What are the current issues impacting scope?
- Are we on schedule? What's late? What is about to start? What is about to finish? What issues are impacting the schedule?
- Are we in budget? Where are we about to go over? What threats are present impacting project costs?
- Who is supposed to be performing activities on the project today? Are all the equipment and materials required available? What issues do we have against resources?
- What is the next milestone? How close are we?
- When the next meeting? What information needs to be communicated?
- What risks are important to monitor today?

This is just a sampling of what the project manager is looking at for the day, or the week. And they are not alone. Though the focus of the information may be different, the project team members and the stakeholders are looking at similar information. Unfortunately, the information that any person is looking at is typically in several locations. In order to get the information, the source needs to be retrieved and scanned.

Using a project management engine as a web application, the information from these multiple sources can be retrieved, compiled, and presented to the user in a single interface. The technique used to make this possible is commonly referred to as a “mashups.” They are web applications designed to allow the user to customize their web interface to allow multiple sources of information to be presented at the same time. With the ability to customize the interface, the user can determine what information needs to be displayed at any given time. Each user could have a different interface to the project, one specifically designed by them to meet their information concerns.

To allow mashups to work, the application developer simply needs to link the source of the information to the application. Some translation may be required to build relationships and associations, but once the information is readable and the criteria are set for how it is read it is ready. At that point, the source becomes a component in the list of components available for the user, like a catalogue. The developer can also create relationships between multiple sets of data; for instance, mapping the project plans to the risk responses to change requests. Each piece of data is found in different sources, but is combined to show in the user's interface of the project.

### **7.3.1 Everyone is a Project Manager (Open Source)**

One of the constraints of the tools used for project management is that they are built for use by project managers. Though nothing is wrong in this, it does alienate other people who are not familiar with project management needs and concerns. In some cases, these people are simply involved to provide resources to the project. Unfortunately, they are being asked to use these tools. And even with project managers, one characteristic starts coming into play: the behaviors of the individual start to change to accommodate the tool. This characteristic always becomes a problem. Extreme adversity towards changing behaviors will result in the tool not being used at all. Modest adversity often leads to not using it to its fullest capabilities or incorrectly.

Though the initial focus is on a project management tool, the same characteristic can be found inside a project. Many projects claim to fail because of lack of teamwork, communication, and understanding. Though these reasons are valid, the more likely source of failed projects is forced behavior changes. When an individual or group is forced to behave contrary to their nature or experience, resistance starts to appear. If the resistance is coming from one or two individuals on a project because they are new to project management, concerns with project failure will not come up, though a risk may be present. But when the entire team are being asked to change their behaviors, the probability of failure increases.

An easy solution to this problem is readily available: encouraged participation in every aspect of the project and the project management tool. The philosophy behind this solution is “open sourcing” the project or tool. How is this done? Starting with the project, it means making the project plan a truly living project. In many project management situations, the individual contributions to the project plan are made and compiled by the project manager. The organization of the project plan is completed by the project manager and reviewed with the project team during a meeting when the plan is distributed. Collaboration is controlled and input to the plan before, during, and after is often filtered. This is the traditional method of project planning. However, what if the project plan is located in a central location that everyone updates? As people update, feedback is provided on specific entries detailing requirements, dependencies, risks, or questions. The critical components of the plan are quickly identified and agreed upon by the team. The project manager might organize the plan, but everyone has a say in its construction. And when the plan is complete, the collaboration continues with change controls in place. Eventually, the project plan will be complete and have a richer quality to it than typically available through traditional means. The trend of open sources has shown that the work is done in far less time than the traditional method because the collaborative component has improved.

The same method can be applied to the project management tool. By allowing users to input their needs into a project management tool, eventually a tool is available for use by all users, not just project managers. The core of the tool or the project does not change, just the behavior changes to match that of the project management team. Such changes to behavior characteristics within an application require extensive budgets to build and maintain customized applications, some of which were never fully used by the users. In the cloud, the application is developed as it's being used as and at a far less investment than traditionally required.

The underlining theme behind open sourcing is collaboration at any time for anything. The theme continues during project execution. Traditionally, when a problem arises during the performance of an activity, the performing party notifies the project manager and the two may attempt to find a resolution. Other team members may get involved at the request of the performing party or the project manager, but the group looking at the problem is still restricted. In a collaborative environment, everyone has the opportunity to be involved. In most cases, the problem will be resolved faster with a better, more creative solution and in project management; the risks involved are minimized to the greatest extent. Some people may claim that the old adage “too many cooks in the kitchen...” but the truth is, in the largest, most prestigious kitchens of the world, there is always more than one cook.

## ***7.4 Treating the Project as Parts, not the Whole (Reuse)***

With mashups, different components could be used to create a user interface. The constructs of these components had to be designed and developed, but once done could be used by anyone. Variables set in the components could be changed by the user to allow for further focus on the information retrieved. With blogs and other RSS feeds, users can identify the parts of the project that they want to be regularly informed. The number of feeds available may reach the hundreds, but the user only wants one or two that suits their needs. The concept of reusing predefined components is prevalent in web design and cloud computing. It provides the development of a particular product to be completed much quicker, at far less cost, and fewer risks.

Especially for organizations that perform projects of a similar nature repeatedly, such as building a house, or application, or a prototype, the concept of reusing project plans is an attractive notion. PMBOK even encourages the use of historical information and templates to create the necessary documentation on a project. But every project is unique, so some care has to be made to ensure that adoption of historical information is not made blindly. This does not mean that reusable capabilities cannot be encouraged.

Through separating a project into its core deliverables, a project can be split into components. Each component will have its basic relationships, requirements, dependencies, and risks defined. These interactions will exist internally to the components as well as the external interactions identified. By having multiple components available, a project team can identify the components that are required for the project and add it. Changes to the component for a specific project can be noted and used as historical data for the next project without changing the core component. Thinking of the project as a product, the ability to reuse components provides a cheaper, quicker product to market than recreation from scratch. By breaking projects down into reusable components, the technique is now more scalable, flexible, and reliable than traditional methods.

How does cloud computing support this endeavor? By using links. Entire components can be added to a project plan with the click of a mouse. In minutes, the majority of a project can be created with all the relationships and pertinent information intact. At this time, cloud computing solutions allow business applications to be built in a matter of minutes using components already available. The concept is the same, using the project as the basis for the application. Referring to the relationships is not restricted to the relationships between information, and also the processes of project management. The different components can include the various tools used by project management, like a change control system, or risk response system, or project monitoring systems, or performance reporting. Normally these systems are not interrelated. In web applications, the business logic can stay where it exists, however the presentation layer can show information, workflow, and indicators to be used actively by the project team and stakeholders.



## **7.5 Testing the Limits (*Portability*)**

Portability is a great advantage to any endeavor. By using the web as the basis for building, monitoring, and executing a project, portability is automatically provided. Using a browser and a connection to the Internet, a user can access a project from any location in the world, at any time. In many situations, the performing party is executing on the project several tasks defined by the project at any given times. Many tasks, especially in manufacturing and construction, require not being readily accessible to the project plan to update the situation. So the performing party waits to the end of the day to update the project plan and any related materials. This is particularly valid when the software used to update the project plan is sitting on the computer at the office. However if the performing party had a hand held device able to connect to the Internet like a phone, they could make an update to the project plan on the spot.

The concept of portability is a simple one. More important are the benefits that appear by using web-based project management. Many project management programs are very expensive and providing a license to every member of a project management team may not be a cost effective solution. Many project managers have filled the gap by taking status updates from members and updating the plan themselves. However, everyone typically has access to a browser allowing them to update the plan as members of the project team. This potentially reduces the administrative work performed by the project manager while also significantly reducing the cost of licensing the tools. In cloud computing solutions, the usage is typically monitored and charged, so a company may have several full-time project members with a couple of dozen part-time members. Traditionally, everyone had one full license despite the usage. With a cloud solution, the usage of the members can be added up and charged. So, inactive members would not prove to be an additional license cost for the company.

## 7.6 Process Review Questions

1. Security is an important consideration when choosing a cloud computing provider.  
Which two processes work together on security policy?
  - a) Configuration Management and Information Security Management
  - b) Availability Management and Information Security Management
  - c) Service Level Management and Information Security Management
  - d) Capacity Management and Information Security Management
2. Cloud computing is more than offsite data storage  
  
T            F
3. Which of the following are examples of Cloud Computing lingo?
  - a) Organization as a Service, Data as a Service, Hardware as a Service, Nothing as a Service
  - b) Computers as a Service, Platform as a Service, Software as a Service, Framework as a Service
  - c) Software as a Service, Data as a Service, Nothing as a Service, Users as a Service
  - d) Framework as a Service, Computers as a Service, Platform as a Service, Organization as a Service
4. Which of the following best describes the changes that occur when cloud computing is utilized?
  - a) The outcome for the end-user does not necessarily change but the delivery strategy is transformed
  - b) The delivery strategy remains the same but the customer accesses the application via an internet connection
  - c) Responsibilities and accountabilities may see potential changes but the service management processes remain the same
  - d) Nothing will change when cloud computing is utilized
5. SLAs would NOT include:
  - a) Scope
  - b) Service availability
  - c) Financial trends
  - d) Service performance

6. Which section sits underneath the core business processes in the objective tree?
- a) IT service organization
  - b) IT technical processes and functions
  - c) IT service management
  - d) Organization
7. Service level management consists of three phases: Budgeting, Development and Delivery.
- T          F
8. Which of the four service management perspectives has traditionally been given the most attention?
- a) Processes
  - b) Organization
  - c) People
  - d) Technology
9. When cloud computing is used, control may be lost over:
- a) Security
  - b) Storage
  - c) Privacy
  - d) All of the above
10. What is a virtual service desk?
- a) The service desk is co-located within the user community it serves
  - b) Service desks are merged into a single location by drawing staff into one or more centralized service desk structures
  - c) The impression of a single, centralized service desk is given when in fact the personnel may be spread or located in any number or type of geographical or structural locations
  - d) Two or more service desks are dispersed geographically to provide 24-hour service

*Answers to all questions can be found in Chapter 7.*

# Answers

## 1.9 Introduction

### ANSWERS

1d, 2b, 3F, 4c

## 3.29 Virtualization Technologies

### ANSWERS

1d, 2a, 3T, 4b

## 7.6 Process Review

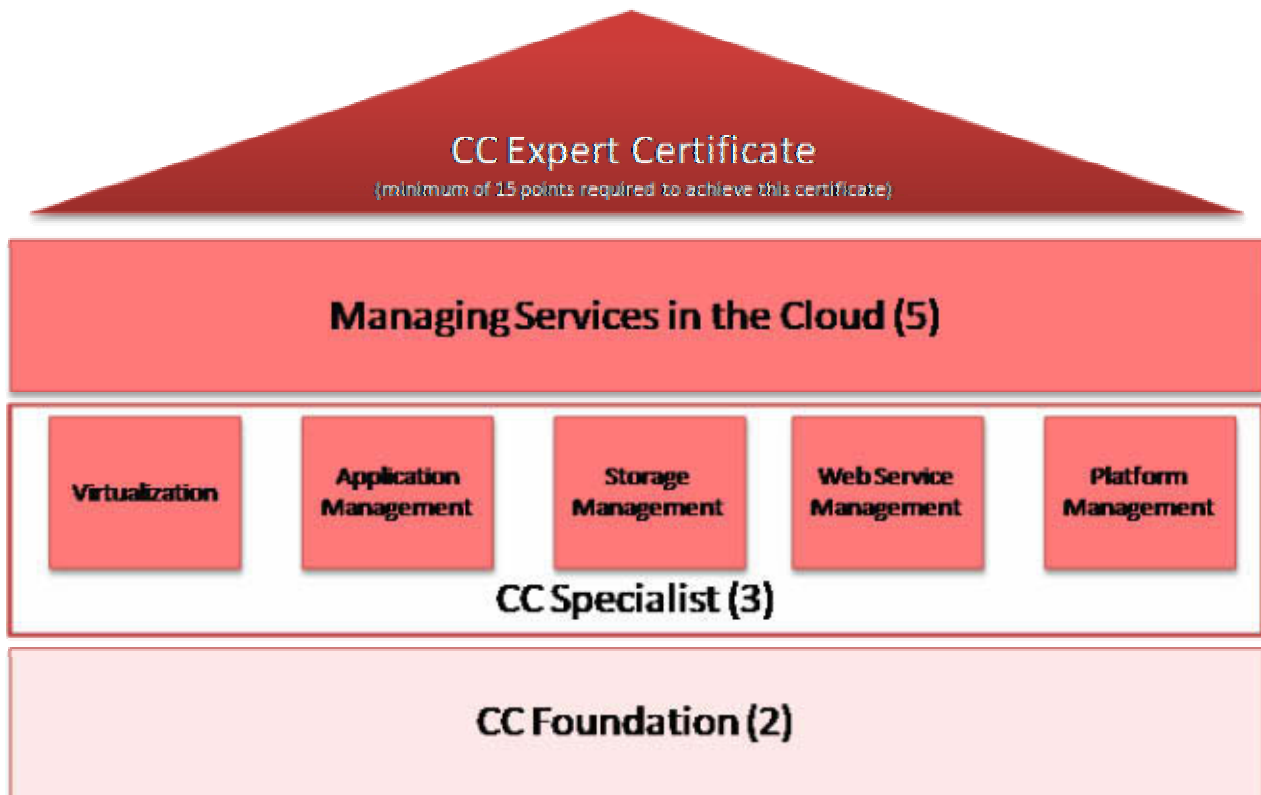
### ANSWERS

1b, 2T, 3a, 4d, 5c, 6a, 7F, 8d, 9d, 10c

# Certification

## 7.7 Cloud Computing Certification Pathways

IT professionals need to know a whole lot more about the various ways of delivering services to the customers and end-users. It is no longer sufficient just to know the differences between Windows-based or Linux-based architecture. These days, most services will utilize some form of Cloud Computing, be it virtualization or SaaS offerings. So with the change in computing and IT Service Delivery comes a whole new series of qualifications and certification. The Cloud Computing Certification Scheme has been created to support the IT Professional who needs to be a 'niche generalist', especially in a rapidly changing area like Cloud Computing.



### How does the certification pathway work?

First, you need to create the foundation – ***The Cloud Computing Foundation Program*** focuses on the fundamentals, general knowledge, terminology and BASIC concepts used in Cloud Computing. This program earns you **2 points** toward your Cloud Computing Expert Certificate.

After this come the Cloud Computing Specialization options. The prerequisite for these programs is that you have the Cloud Computing Foundation certificate.

We appreciate that you don't need to know everything about each area of the IT organization so this is where the programs become more specialised:

- **Virtualization**
- Application Management (including SaaS)
- Storage Management
- Web Service Management (Hosted solutions)
- Platform Management (Platform as a Service)

Each program in this series earns **3 points** towards the Cloud Computing Expert Certificate.

The next level is 'Managing Services in the Cloud' and this program is specifically aimed at Service Managers and IT Service Delivery Managers who wish to add Cloud Computing as an option in their organization's delivery model. The program is worth 5 points.

### **How do I achieve my Expert level?**

You must have a minimum of 16 points to achieve your expert certification. There are 2 pathways to follow:

- a) **Technical Specialist Pathway:** Cloud Computing Foundation Program + 5 x Cloud Computing Specialist Programs

**Total:** 17 Points

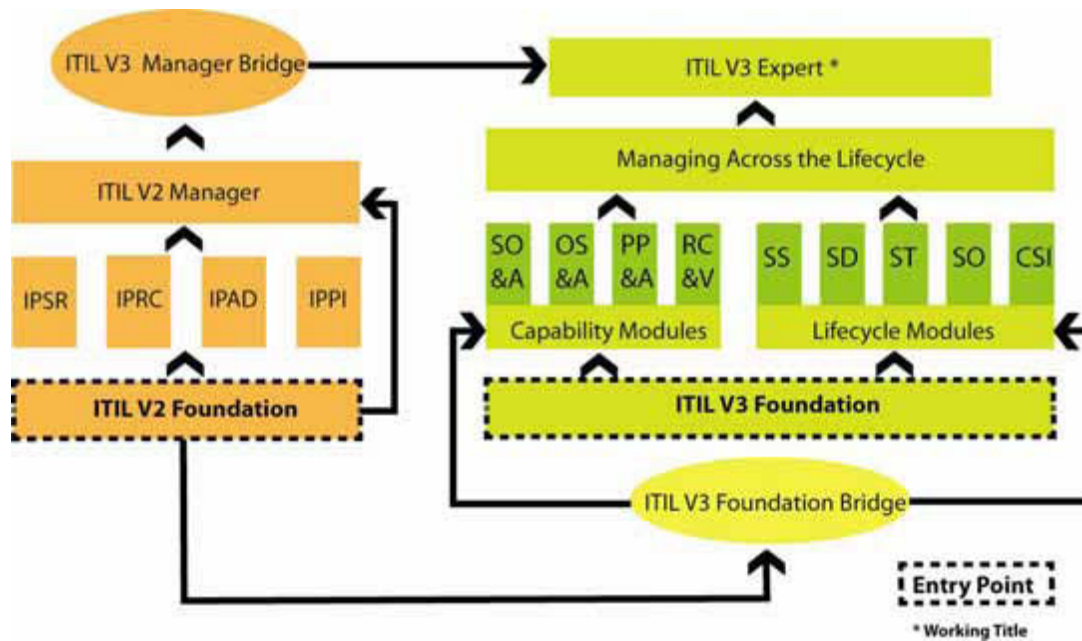
- b) **Manager Pathway:** Cloud Computing Foundation Program + 3 x Cloud Computing Specialist Programs + Managing Services in the Cloud Program

**Total:** 16 Points

Each course provides preparation for the exam and successful candidates receive a certificate results sheet and a lapel pin.

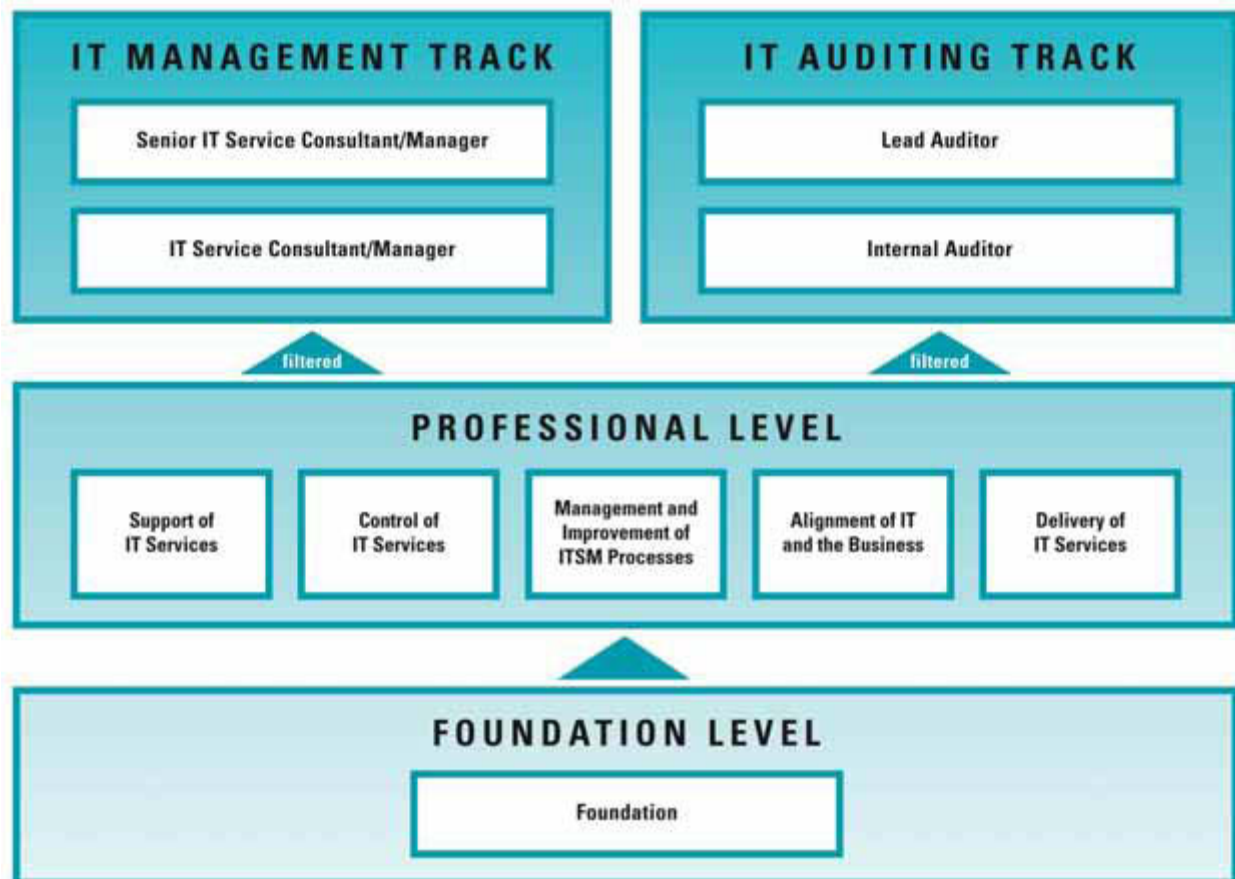
## 7.8 ITIL® Certification Pathways

There are many pathway options that are available to you once you have acquired your ITIL® Foundation Certification. Below illustrates the possible pathways that are available to you. Currently it is intended that the highest certification is the ITIL® V3 Expert, considered to be equal to that of Diploma Status.



## 7.9 ISO/IEC 20000 Certification Pathways

ISO/IEC 20000 Standard is becoming a basic requirement for IT Service providers and is fast becoming the most recognized symbol of quality regarding IT Service Management processes. ISO/IEC 20000 programs aim to assist IT professionals to master and understand the standard and the issues relating to earning actual standard compliance.



For more information on certification and available programs please visit our website <http://www.artofservice.com.au>



## 8 Virtualization Specialist Exam Tips

### Exam Details

- 20 multiple-choice questions
- The correct answer is only one of the four
- 30 minutes duration
- 14 out of 20 is a pass (60%)
- Closed-book
- No notes

### Practical Suggestions

- Read the question CAREFULLY.
- At this level of exam the obvious answer is often the correct answer (*if you have read the question carefully!!*).
- Beware of being misled by the preliminary text for the question.
- If you think there should be another choice that would be the right answer, then you have to choose the “most right”.
- Use strategies such as “*What comes first?*” or “*What doesn’t belong?*” to help with the more difficult questions.

### Organising your Exam

The Art of Service facilitates the Virtualization Specialist exam. Please contact us on 1300 13 44 99 or email [exams@theartofservice.com](mailto:exams@theartofservice.com) to arrange your examination.

Make sure that you prepare adequately in the lead up to your exam by reviewing your notes, reading any available material and attempting the sample exams.

We wish you luck in your exam and future Cloud Computing & Virtualization career!

## 9 References

Carey, W.P (2008). *Cloud Computing: The Evolution of Software-as-a-Service*. Found at: <http://knowledge.wpcarey.asu.edu/article.cfm?articleid=1614>

Foley, J (2008). *20 Cloud Computing Startups You Should Know*. Found at: [www.informationweek.com/news/software/hosted/showArticle.jhtml?articleID=210602537](http://www.informationweek.com/news/software/hosted/showArticle.jhtml?articleID=210602537)

Higginbotham, S (2008). *10 Reasons Enterprises Aren't Ready to Trust the Cloud*. Found at: <http://gigaom.com/2008/07/01/10-reasons-enterprises-arent-ready-to-trust-the-cloud/>

Jones, T (2008). *Cloud Computing with Linux*. Found at: <http://www.ibm.com/developerworks/linux/library/l-cloud-computing/index.html>

King, R (2008). *How Cloud Computing is Changing the World*. Found at: [http://www.businessweek.com/technology/content/aug2008/tc2008082\\_445669\\_page\\_3.htm](http://www.businessweek.com/technology/content/aug2008/tc2008082_445669_page_3.htm)

MacVittie, L (2008). *The Three "Itys" of Cloud Computing*. Found at: <http://devcentral.f5.com/weblogs/macvittie/archive/2008/09/17/3622.aspx>

Mathson, M. *Brief Introduction to Cloud Computing*.

Scanlon, J.H & Wieners, B (1999). *The Internet Cloud*. <http://www.thestandard.com/article/0,1902,5466,00.html>

SYS-CON Media (2008). *What's the difference between Cloud Computing and SaaS?* Found at: <http://web2.sys-con.com/node/612033>

The Art of Service (2008) *CMDB and Configuration Management Creation and Maintenance Guide*, Brisbane, The Art of Service

The Art of Service (2008) *ISO/IEC 20000 Foundation Certification Kit*, Brisbane, The Art of Service

The Art of Service (2008) *Risk Management Guide*, Brisbane, The Art of Service

The Art of Service (2008) *IT Governance, Metrics and Measurements and Benchmarking Workbook*, Brisbane, The Art of Service

The Art of Service (2007) *ITIL Factsheets*, Brisbane, The Art of Service

The Art of Service (2008). *ITIL® V3 Foundation Complete Certification Kit*. United Kingdom, Emereo Pty Ltd.

Wikipedia. *Cloud Computing*. Found at: [http://en.wikipedia.org/wiki/Cloud\\_computing](http://en.wikipedia.org/wiki/Cloud_computing)

Wordpress (2008). *Clouds: Dark and Threatening or the Nice, Fluffy Ones?* Found at: <http://itsuperhero.wordpress.com/2008/10/01/clouds-dark-and-threatening-or-the-nice-fluffy-ones/>

<http://technet.microsoft.com/en-us/library/bb963714.aspx>

<http://www.virtualizationworks.com/Virtualization-History.asp>

<http://www.eweek.com/c/a/Virtualization/VMware-Cisco-Team-to-Enhance-Data-Center-Virtualization/>

<http://windowsitpro.com/article/articleid/98955/5-virtualization-companies-to-watch.html>

<http://www.vmware.com/virtualization/>

[http://h71028.www7.hp.com/enewsletter/cache/564109-0-0-225-121.html&pagesite=SMB\\_LIVE\\_OOV](http://h71028.www7.hp.com/enewsletter/cache/564109-0-0-225-121.html&pagesite=SMB_LIVE_OOV)

<https://help.ubuntu.com/8.04/serverguide/C/virtualization.html>

<http://www.ubuntu.com/products/whatisubuntu/serveredition/technologies/virtualization>

<http://www.microsoft.com/windowsserver2008/en/us/hyperv.aspx>

<http://www.networkworld.com/reviews/2008/090108-test-virtualization.html?page=1>

<http://www.redhat.com/rhel/virtualization/>

<http://www.techsoup.org/learningcenter/software/page4826.cfm>

<http://www.thinstall.com/>

<http://in.sys-con.com/node/689946>

<http://blogs.zdnet.com/virtualization/?p=311>

<http://en.wikipedia.org/wiki/Wiki>

[www.artofservice.com.au](http://www.artofservice.com.au)

[www.theartofservice.com](http://www.theartofservice.com)

[www.theartofservice.org](http://www.theartofservice.org)

**INDEX\***

32-bit 32-3, 50, 54  
 64-bit 50, 54

**A**

ability 21, 24, 28-9, 33, 42, 45, 61, 68, 75, 92, 125, 127-9, 158, 167-8, 171, 173  
 Ability to Separate Applications 14  
 acceptance criteria, Service 112  
 access end-users 34  
 Accomplishing Virtualization 76  
 administrators 10-14, 27-8, 32-4, 37-8, 40-3, 45, 52-3, 55, 58-9, 63, 65, 71-2, 74-5, 80, 89, 104-5 [1]  
 advancements 99-100  
 agents 162  
 agreements 5, 101, 103  
 allocations 18, 165  
 Altiris 27, 29  
 amount of hardware emulation 39  
 appliance 162, 164  
 application change 113  
 application components 109  
 application delivery 21  
 application developer 171  
 application development 167  
 application layer 159  
 application level 99  
 application-level virtualization technologies 26  
 application processing 18, 20  
 application productivity 95  
 application servers 163  
 application software 129  
 application stop 129  
 application versions 72  
 application virtualization 4, 29, 41-2, 46, 63, 75-6, 87  
 application virtualization program 47, 54, 57-8, 62, 65, 88  
 application virtualization software 46, 63, 99  
 application virtualization solutions 63, 65, 87  
 applications 12-18, 21-4, 27-30, 33-4, 40-3, 45-9, 56-8, 62-6, 79-81, 83-5, 87-90, 92-3, 96-9, 158-62, 166-8, 171-3 [8]  
   associated 8  
   based 64  
   customized 172  
   desktop 167  
   destructive 58  
   downloadable virtual machine benchmark 53  
   downloading Altiris's SVS 29  
   extra 72  
   financial 137  
   freeware 30  
   group 45, 75  
   high performance 49  
   hosted 46  
   install 29, 34, 72  
   intended 165  
   keep 72  
   legacy 57  
   malicious 35, 62  
   multiple 11, 40, 63  
   multitasking 40  
   non-native operating system 42, 75  
   non-Oracle 74  
   normal 62

- open source 47
- open source QEMU processor emulator 50
- packaged 63
- packaged 63
- packaging 57
- pre-installed 26
- proprietary 57
- real-time 85
- required 35
- running 87
- running Dot Net 64
- updated 34
- vendor s 97, 128
- virtual 41, 57
- virtual server 28
- virtualized 43, 54
- virtualizes 47, 54
- applications code 40
- applications stops 98
- Applying Cloud Computing to Project 167
- approach, push 114
- Art of Service 182-4
- aspx 183-4
- assemblies 111
- assessment 133, 137-9
- assets, virtualized 52
- associations 170-1
- authorization 133, 136, 138-9
  - local 137-8
- automation 112, 114, 160, 165
- availability 17, 67, 107, 125, 128, 140, 142, 149, 161, 164-6
  - high 129, 164

## **B**

- backup solutions 69, 89
  - efficient 43, 71
- balance 83, 128, 130, 132, 165
- bandwidth 95-6
- bandwidth requirements 96
- Banking 110
- bare metal environment 21, 23, 84
- bare metal solution 48, 76, 84, 98
- Bare metal virtualization applications 46
- bare metal virtualization solutions 30, 53, 81, 84-5, 87-8, 99
- based operating systems 11, 50-1, 53, 56
- baseline 119, 138, 145
- Big Bang 113-14
- black boxx201f 101
- block virtualization 161
  - functions of 17, 161
- blocks 15, 17, 158, 160-1
  - logical 15, 160
- blogging 6, 169-70
- blogs 169-70, 173
- blogs.zdnet.com/virtualization 184
- Board 137-8
- Brisbane 183-4
- browse 27, 72
- browser 168, 174
- budgets 87, 106, 170, 172
  - Business 156
- business activities 123

- business applications 173
- business area 143
- business benefit 128
- Business Case 145
- business change 109
- business change processes 134
- business disruption 108
- business enterprises 43
- business environment 38, 92, 97
- Business Executive Board 138
- business goals 109, 142, 159
- business impact 128
- business logic 168, 173
- business management 137
- business manages 100
- business office 13
- business operations 128
- business/organization 122
- business process 130
- business purchases 103
- business risk 130
- Business Value 6
- Business Value for Virtualization 166
- businesses 10, 12-14, 31, 40-1, 43-4, 58-60, 67, 70-2, 76-7, 82-107, 128-32, 139, 141, 145, 149-50, 158-9 [9]
  - large 8, 38
  - time 83
- businesses infrastructure 35
- businesses migrate 70, 83
- business s 58, 71, 76, 87, 104
- business s availability 61
- business s flexibility 40
- business s management strategy 101
- business s office 106
- business s power 71
- business s requirements 84
- business s storage solution 69
- business s transition 80

## C

- CAB (Change Advisory Board) 132, 136-8, 141-2
- CAB meeting 142
- caching 162, 164-5
- capacity 9, 15, 17, 77, 101, 146, 160-1, 165-6
- cent uptime 95
- centralized Service Desk structure 152
- centralized service desk structures 176
- certification 6, 177, 181
- Certification Pathways 7, 180-1
- Change Advisory Board, *see* CAB
- change authority 136, 138
- Change Management 5-6, 107-8, 112, 116, 124, 129-30, 132-4, 136-7, 139-44
- Change Management process 130-1, 133, 135, 142
- Change Manager 132, 138-9, 141-2
- Change Models 6, 132, 136-7, 140
- change proposal 136-7, 145
- change requests 135, 137, 171
- Change Schedule 134, 138-9
- CI, *see* Configuration Item
- CI (configuration items) 130, 134, 140, 143, 145
- classification 136-7
- CLI (command line interface) 34

- Client-Hosted Desktop Virtualization 37
- client-server model 11
- clients 14, 16, 58, 159-61
- cloud 18, 21, 167, 172, 183-4
- Cloud Application 21
- cloud computing 4, 6, 8, 17-18, 20-1, 158, 167-8, 173, 175-9, 183-4
- Cloud Computing & Virtualization 182
- Cloud Computing Expert Certificate 178
- Cloud Computing Foundation Program 178-9
- cloud computing solutions 173-4
- cloud computing solutions exploit 167
- Cloud Computing Specialist Programs 179
- cloud computing support 173
- Cloud Service 21-2
- CMS (Configuration Management System) 108, 127-8, 130, 136
- co-exist 63, 65, 71
- code 56-7
  - application s source 51
- collaboration 172
- command line interface (CLI) 34
- communication, general 149, 156
- companies 18, 26, 28-9, 31, 54, 60, 70, 159, 163, 166, 174
- compatibility issues 12, 163
- completion 124, 140
- complexity 22, 59, 97, 109, 119, 129-30, 132, 137, 156, 159-60
- components 33, 110-11, 116, 119, 138, 159-60, 164-5, 168, 171-3
  - required 119
  - service/release 122
  - updated service 114
- computer infrastructure 21-2
- computer technology 11-12, 99
- computers 18, 21, 35, 68, 158, 162, 169, 174-5
  - single 8
- computing 4, 17-19, 158, 177, 184
  - grid 4, 18, 22
  - utility 4, 19
- concept 11, 29, 36, 40, 96, 114, 127, 158, 167, 173-4
- configuration 58-9, 90, 128, 166
- Configuration Item (CI) 130, 134, 140, 145
- configuration items, *see* CIs
- Configuration Management 143-4
  - Service Asset & 19
- Configuration Management System, *see* CMS
- configure 34, 59
- conflicts 14, 29, 40-2, 57, 63, 65, 75, 90
- confusion 91-2, 107
- conjunction 52, 64, 108, 131
- connect 11, 68, 95, 174
- connections 18, 107, 168-9, 174
- Consolidating Service Resources 88
- constraints, service delivery 125
- construction 10, 168, 172, 174
- contact 77, 98, 129, 146, 149-50, 152, 162, 182
- containers 36-7, 45-6, 56, 63, 75, 101
  - single 101
- context 126
- Continual Service Improvement 124
- contracts 5, 101, 103
- control 18, 34, 40, 48, 55, 57, 104, 107, 129, 144, 150, 159, 162-3, 176
- conversion 76, 83, 161
- conversion tools 78, 82-3
- Cooperative Virtualization 36



- Coopvirt 36
- copy 14, 37, 41, 54, 57, 63, 65, 74, 89, 93, 102
- core business processes 176
- cost-saving benefits 103
- costs 13, 22, 40-1, 77, 81, 83-4, 94-5, 103, 105, 124-5, 130, 137, 139, 146, 151-2, 173-4 [5]
  - licensing 94
- customers 21, 98, 108-9, 114, 120, 124-5, 130, 139-41, 143, 156, 177
- customer s business operation 138
- customize 37, 171

**D**

- damage 73, 88, 91-2
- data, amount of 16, 160
- Data, Information, Knowledge, Wisdom (DIKW) 126
- data, transfer of 27
- data centre managers 59-60
- data migration 166
- data storage 17, 59, 158
- data storage virtualization 4, 58-60
- data storage virtualization technologies 59-60
- Data Storage Virtualization 58
- databases 16, 127, 159, 161, 168
  - multiple 168-9
- day business operations 13
- deactivate 27-8
- decision 81, 106, 125, 127, 136, 147
- Defining Virtualization 6, 159
- Definitive Media Library, *see* DML
- delivery 21-2, 115, 120, 130, 176
  - Service 123
- delivery of computer infrastructure 21-2
- delivery strategy 175
- departments 103-5
- dependencies 170, 172-3
- deploy 34, 90, 93, 99, 108-9
- Deploy Service Management Capability 122
- deploy software applications 42
- deployment 22, 75, 92, 107-9, 111-12, 114, 116-20, 122-4, 134, 140
  - rapid server 28
- deployment activities 5, 115, 117
- deployment management 19, 139-40
- deployment plans 108, 112, 115-16, 120, 122
- deployment process 112, 115-17
- Desktop Edition 26, 30
- desktop virtualization technologies 37
- devices 16, 159-60, 174
  - networking 96
- DIKW (Data, Information, Knowledge, Wisdom) 126
- disk 15-16, 24, 160-1, 165
  - magnetic 15, 160
- disk virtualization 15, 160
- disruptions 14, 128, 130, 143, 149, 157
- distribute 107-9
- distribution 107, 109, 112, 168
- DLLs 33, 41, 57, 62-3
- DML (Definitive Media Library) 111
- documentation 112, 118, 133, 142, 173
- documentation templates 118
- domain 100-1, 141
- domain names 45
- dot net framework 64-5
- Dot net framework virtualization 65

download 30-1, 53, 75, 114  
downtime 10, 14, 61, 76, 87, 98, 166  
    schedule 95  
drive, virtual 16, 160-1  
drivers 41, 168

## **E**

Early Life Support 116, 122-3  
ECAB (Emergency CAB) 136, 138, 142  
edit 41, 53, 55  
efficient virtualization technology 84  
effort 70, 119, 128, 164  
email exams@theartofservice.com 182  
EMC 60  
Emergency CAB (ECAB) 136, 138, 142  
Emergency Change 133, 142  
employees 92, 104-5, 133  
emulate 16, 36, 44, 55, 86, 158, 160  
    virtual machine software 22  
emulate hardware 39, 41, 45  
end-user community 117, 143, 149  
end-user desktops 37-8, 78  
end-users 21, 59, 68-9, 143, 175, 177  
engines 163-4  
enterprise level virtualization product 30  
enterprise network storage architectures 60  
entities, separate 8, 10-11, 13, 20, 74, 107  
entries, registry 29, 33, 41, 57, 62-3  
environments 8, 10, 14, 16, 18, 27-8, 31, 43, 46, 63, 79, 98, 116-17, 129-30, 163, 165-7 [7]  
    business s 61  
    container 45  
    hosting 87  
    isolated 29, 42, 45, 47, 58, 62, 65  
    live 107, 109, 111, 119, 140  
    server-based 68  
    users 114  
errors 12, 82, 114, 130, 145, 165  
Establishing data 125  
evaluation 116, 124, 131, 134, 136, 140, 142, 146  
events 74, 139, 164  
Evolution of Software-as-a-Service 183  
exam 179, 182  
expenditures 106

## **F**

factors, external 100  
failure 71, 89, 107, 133, 143, 152, 159, 171  
Fedora 5, 66  
file 13, 15-17, 57, 67, 71, 77, 160-1  
file systems 15-16, 159-63, 166  
file systems virtualization 16  
file virtualization 16, 161  
Financial Management 5, 105  
Financial Management and Virtualization 106  
firewalls 72, 92  
fluctuations 16, 160  
Follow the Sun Service Desk structure 154  
Framework 175  
full virtualization 8, 22-3, 25, 28, 32, 44-6, 63  
    hardware-assisted 36  
full virtualization solution 26, 33, 45-6, 63, 87-8  
full virtualization technologies 43

functionality 10, 13, 33, 160  
 functions 108, 150, 159, 162-3, 166, 176  
     advanced virtualization 163

**G**

GB 77  
 Google 18  
 graphical user interface (GUI) 26, 34, 47, 51-2, 57  
 groups 12, 45, 101, 106, 117, 130, 133, 136, 138, 140, 171-2  
 growth, restricting business 132  
 guest 26, 30-2, 44, 46, 51-2, 54-5, 74, 99  
 guest machine 35-6  
 guest OS 22, 24, 28, 33, 44  
 GUI, *see* graphical user interface

**H**

hard drive 8, 13, 25, 43, 66, 68-9, 77-8  
     physical 23-4  
 hard drive partitions 41  
 hardware 8, 22, 26, 30-3, 35-6, 38-9, 44-6, 54-6, 58-60, 63-4, 74-5, 77, 79-84, 86-7, 94-5, 102 [7]  
     amount of 102-3  
     physical machine s 86  
     server-based 49  
     server s 30, 84  
     supported 54  
     virtual machine s 24  
 hardware assistance 33, 88  
 hardware-assisted virtualization 22, 36, 38-9, 44, 46, 54, 60-1  
 Hardware-assisted virtualization solutions and Para-virtualization solutions 86  
 hardware-assisted virtualization technologies 54-5, 85-6  
     adopted 39  
 hardware-based virtualization technology 33  
 hardware costs 10, 40, 94  
 hardware emulation 8, 39, 46  
     software-based 61  
 hardware faults 10, 14, 20  
 hardware features 73, 79  
 hardware formats 86  
 hardware platforms 52, 86  
 hardware requirements 77, 81  
 hardware resources 8-9, 11-12, 20, 25, 31, 39, 63  
     mainframe machine s 11  
     physical 39  
 hardware resources split 12  
 hardware solutions, best 80  
 hardware support 44  
 hardware virtualization 4, 12, 26, 29, 38, 75  
 hardware virtualization technologies 26, 39, 41  
 heads 15, 160  
 Hierarchical Storage Management (HSM) 16, 161  
 Hiring Service Desk 155  
 History of Virtualization 11  
 HISTORY of VIRTUALIZATION 4  
 host 21, 23, 28, 44-6, 52, 55, 58, 61-2, 72, 74, 76-8, 84, 86-9, 93-4, 99, 162-6 [1]  
 Host-based virtualization 162  
 host environment 165  
 host layer 161, 165  
 host machine 42, 74, 78, 92, 94, 96  
     physical 72, 74, 91  
 host OS 22-4, 29-30, 33, 37, 63  
 host OS level 23  
 host server 162

- hosted environment 21, 23
- hosted virtualization solutions 51, 53-4, 82, 84-9, 99
- hosting 23, 45, 63, 91, 95, 100
- HP 47, 60
- HP Integrity Virtual Machines 47
- HR (Human Resources) 133, 155
- HSM (Hierarchical Storage Management) 16, 161
- Human Resources (HR) 133, 155
- Hyper-Threading 40
- hyperlinks 6, 168, 170

**I**

- IBM 11, 60
- IDEs 68, 167
- implementation 117, 122, 138-40, 143, 150, 153, 158, 161-3
- implementation process 80, 163
- impression 149, 153, 176
- Improved business flexibility 9, 20
- Improved Business Flexibility 10
- Improved customer service perception 150
- Improved quality of service 154
- improvements 9, 108, 130, 142, 150, 164-5
- Improving Availability Using Virtualization 6, 164
- in-band approaches 163
- Inc 50, 53, 55
- Incident Management 5, 128
- Incident Management and Virtualization 128
- incidents 111, 122-3, 133, 143, 149, 152, 155-7
- Independent Software Environments 167
- infiltration 70
- information 13, 98, 100-1, 107, 119, 125-8, 132, 137, 158, 162, 168-71, 173, 181
  - historical 167, 169, 173
- information security 107
- Information Security Management 5, 175
- Information Security Management and Virtualization 107
- infrastructure 18, 32, 40, 61, 76, 82, 85, 87-8, 92, 96, 104, 109, 115, 128, 139, 156 [1]
  - business s 77, 101, 103
- Infrastructure-as-a-Service 22
- initiator 136-7, 140
- inputs 34, 112, 134, 172
- install 21, 46, 57, 62-3, 82, 89-90
- install software applications 27
- installation 54, 65, 75-6, 115, 119
- installing software applications 57
- instances 78, 80, 89, 102, 143, 159, 164, 167, 169-71
- instructions 35-6, 44, 75
- Intel 26, 39, 41, 100
- Intel and AMDs hardware virtualization technologies 66
- Intel processors 53-4
- Intel virtualization 31, 61
- Intel Virtualization 4, 31
- Intel virtualization technology 60-1
- Intel Virtualization Technology 5, 60
- Intel VT 26, 31-3, 36, 39, 41, 60-1, 66
- Intel VT technology 31, 61
- Intel s hardware virtualization technology 39
- interact 39, 46, 84, 87
- interfaces 5-6, 30, 68, 79, 96, 109, 112, 131, 134, 168, 171
  - common 27
- internet 17-18, 24, 72, 114, 153, 168, 174
- IP address 90
- iSCSI 67-8

- ISCSI 68
- ISO/IEC 7, 181, 183
- isolation 32, 57, 87
  - complete 42, 87-8
- IT Service Continuity Management (ITSCM) 19, 147
- Itanium 47
- Itanium processor 47
- items 122, 124, 157, 170
- ITIL 7, 132, 180, 184
- ITSCM (IT Service Continuity Management) 19, 147

**J**

- job position 168
- judgments 127

**K**

- keep track 59, 74, 91-3, 102
- kernel 26-7, 36, 40-1, 44-6, 66
  - operating system s 44-5, 61, 75
  - single 36, 40, 45, 84, 86, 88
- Kernel-based virtual machine 26
- Kernel-based Virtual Machine 81
- Kernel-Virtual Machine 36
- Key Performance Indicators (KPIs) 6, 142, 157
- KEY PERFORMANCE INDICATORS (KPIS) 6, 142, 157
- knowledge 18, 98, 105, 108, 123, 125-7, 129, 152, 163, 178
- Knowledge Management 5, 125-6
- Known Error 110, 124, 154
- KPIs, *see* Key Performance Indicators
- KPIS, *see* KEY PERFORMANCE INDICATORS
- KVM 26, 36, 66, 81

**L**

- language 56, 150, 152
- layer 59, 101, 159, 161, 163-4
  - block aggregation 159
  - virtual 27, 29
- LBA (logical block addressing) 15, 160
- levels 69-70, 87, 98-9, 111, 120, 128-9, 137-8, 150, 159, 161, 169, 179, 182
  - appropriate 137, 139, 142
  - file system 69-70
  - service desk staffing 156
- licence 30, 102, 119, 124
- licenses 89-90, 94, 102, 117, 174
- licensing 5, 89, 94, 101-2, 174
- links 168, 170-1, 173
- Linux 11-12, 26, 30, 35, 47-8, 50-6, 94, 183
- Linux-based application 48, 51
- Linux-based distributions 66, 99
- Linux-based operating systems 26, 32, 36, 49, 56, 81
- Linux-based technologies 36-7
- Linux-based virtual machine 99
- Linux-based virtualization management application 57
- Linux-based virtualization management program 48
- Linux distributions 36, 50, 55-6
- Linux users 56
- Linux Virtualization 4, 35
- Live PC Engine 49
- LivePC virtualization technology 49
- load 42-3, 53, 66
- Local Service Desk 151
- local Service Desk structure 151

- Local Service Desk structure 151
- locations 25, 42, 70, 90, 92, 97, 114, 125, 128, 148, 150, 152, 167, 169-70, 172, 174
  - physical 8, 10, 16, 59, 158, 161
  - target 113-14
  - virtual 27
- Log files 74
  - virtual machine s 74
- logical block addressing (LBA) 15, 160
- logical volume managers (LVMs) 162
- LUNs 158, 162-3, 165-6
  - single 162
- LVMs (logical volume managers) 162

## **M**

- Mac 42, 49, 51
- Mac application 50
- Mac-based full virtualization 50
- Mac-based virtual machines 48
- Mac-based virtualized application 54
- Mac-on-Mac (MoM) 49
- Mac OS X 50-2, 55
- machines 8, 10, 12-14, 21, 23-4, 26-8, 35, 37, 41-4, 56-8, 61-6, 70-3, 79-80, 82, 90-4, 102-3 [7]
  - entire virtual 13-14, 23, 37, 83, 107
  - five virtual 77, 94-5, 100
  - managing virtual 55, 104
  - networked 22
  - new 10, 74-5, 80, 90
  - single 27, 31, 40-1, 43, 45, 65, 68, 71, 96, 103
  - software-based 91
  - thin client 38, 52
  - virtualized 23, 33, 38, 48, 77
- machines host 57
- machine s 31, 64, 66, 78, 95
  - physical 77, 95
- machine s performance 100
- machine s storage capacity 103
- Macintosh users 53
- management 9, 19, 51, 59, 61, 69, 83-4, 87, 101, 112, 125, 137-8, 140, 144, 150, 160 [2]
  - service level 176
  - Application 178
- management application 52
  - popular 55
- management costs 12, 20, 61, 83, 85, 94
- management solution 87, 101
- management tools 83
- Managing Services 179
- mapping, address space 158
- mashups 170-1, 173
- meeting, service level management review 141
- members 141-2, 174
- memory 17, 24, 30, 73, 102, 161
  - application accesses 64
  - virtual 8, 25
- meta-data 158
- metrics 101, 108, 142-3, 157, 183
- Microsoft 26, 31, 33, 48-9, 64, 83, 99
- Microsoft Application Virtualization 33
- Microsoft Virtual 26, 37, 48-9
- Microsoft Virtual Server 49
- Microsoft Virtualization Solution 33
- Microsoft Windows 56-7
- Microsoft s 33-4, 48-9

- Microsoft s Management Console (MMC) 34
- Microsoft s Softgrid application 47
- Microsoft s virtualization solution 26
- Microsoft s Virtualization Support Policy 99
- migrate 12, 63, 70, 87, 100
- migrating 5, 76-7, 80, 83-4, 94, 96, 101, 103, 129
- migration 5, 23-4, 76, 78-9, 82, 103, 163, 166
- migration process 76, 78-80, 82-3
- Minor Release 110, 115
- MMC (Microsoft s Management Console) 34
- mode, privileged 35
- model, client/server 97, 128
- modern day businesses 14
- modern day virtualization technologies 96
- MojoPac 49, 87
- MoM (Mac-on-Mac) 49
- money 10, 13, 38, 59-60, 71, 77, 81, 83, 85, 94, 100, 102-3
  - saving 94, 103
- monitoring 74, 99-100, 104, 157, 174
- Most project management solutions 169
- Most Service Desk Managers 155
- Most virtualization approaches 162
- Most Windows-based applications 54
- multiple computers, role of 8
- multiple virtual machines 8, 14, 26, 31, 37, 40, 47, 52, 85-6, 88, 94-6, 102
- multitasking 40

## N

- naming conventions 78, 119
- NAS 67
- NAS devices 67
- NAS technology 67
- Net framework virtualization 64
- Net Framework Virtualization 5, 64
- network 16, 22, 27, 35, 67-8, 72-3, 78-9, 82, 91, 95-6, 130, 159-60, 162, 165-6
  - business s 87
  - organization s 43
- Network-attached Storage 67
- network capability 15
- network connections 81, 101
- network interface card 27, 95
- network interface cards, multiple 96
- network parameters 78
- network resources 39, 159
- network stack 159, 162
- Networking costs 13, 106
- NIC 95-6
- number 10, 15, 30-1, 43, 45-51, 54-61, 64-5, 78-9, 89-90, 102, 104-5, 129, 152-3, 156-7, 167, 169-70 [14]
  - large 16, 22, 42, 61, 160

## O

- objectives 5-6, 9, 108, 130, 141, 149-50
- Objectives of Virtualization 9
- OBJECTIVES of VIRTUALIZATION 4
- one-to-many environment 18, 20
- open source virtualization technologies 31
- operating environments 48, 50-1, 53-4, 57, 62, 64, 66, 74, 81, 84, 87-9
- operating system environment 54, 56, 65, 69, 73, 84, 87
- operating system level 93, 99
- operating system virtualization solutions 81
- operating systems kernel 26, 36, 40

- Oracle 74
- Oracle Virtualization 5, 74
- Oracle VM 74
- Oracle VM templates 74
- Organizational Considerations 5, 103
- organizations 9-10, 12, 41-2, 58-61, 68-9, 77-8, 80-1, 87, 105-7, 124-8, 130, 132-3, 149-50, 155-6, 172-3, 175-6 [17]
  - large 58, 112
- organization's service costs 109
- organizations , global 153-4
- organization s 68, 119, 143
- organization s ability 43
- organization s capabilities 108
- organization s release policy 109
- OS-level virtualization 44
- Os Virtualization 4
- OS virtualization 45-6, 83-4, 86
- OS Virtualization 44
- OS virtualization solution controls 86
- OS virtualization solutions 45, 85-6, 88
- OS virtualization sounds 45
- OS virtualization technology 88
- out-of-band virtualization forces 164
- outsourcing 6, 157
- overarching strategy 115
- overheads 26, 39, 45
- overlook 71, 74, 94, 96

## P

- page, single 168
- Para-virtualization 35-6, 44, 61
- Para-Virtualization 4, 32, 44
- Para-virtualization and Hardware virtualization 38
- Para-virtualization solutions 36, 86
- Para-virtualization technologies 85
- Parallels Desktops 50
- parties 125, 135, 168-9
- partition 22-4, 93
- party, performing 172, 174
- Parvirtualization 61
- pathways 179-80
- PC, virtual 48-9
- PearPC 50
- performance 5, 16-17, 31, 53, 55, 85, 93, 95, 99-101, 157, 161-4, 172-3
  - high 31, 85
  - near-native 84-5
- performance management tools 100
- permalinks 170
- person 54, 92, 138, 142, 146, 167-8, 170
- personnel, unauthorized 93
- perspectives 130, 169
- Phased Approach 113
- physical environment 12, 76, 86, 90-1, 95, 98, 129
  - normal 96-7, 128
- physical hardware 13, 21, 23-4, 35, 44, 46, 64, 69, 85
- physical infrastructure 70, 76, 86, 92
- physical machine logs 74
- physical machines 8, 10, 13-14, 22-4, 27-9, 31-3, 37, 40, 45, 61-2, 71-3, 76-80, 86-7, 89-96, 98-100, 104 [8]
  - five 94
  - normal 71-2
  - securing 74



- physical machines guest OS 56
- physical machine's host 73
- physical machines log entries 74
- physical machines resources 57
- physical machine's BIOS 39
- physical machine's host 73
- physical server 9, 23, 45, 71, 77, 85-6, 91, 94-5, 100, 102
  - single 9, 13, 77, 94
- physical server's ability 94
- physical space 13, 59, 106, 165
- pilots 116-17, 119-20, 143
- planning 6, 76, 78, 80, 113, 115, 120, 131, 164-5
- planning stage 76, 80-1, 83, 96
  - initial 76, 81, 93-4
- plans 80, 112, 115-16, 118, 120, 122, 144, 147, 172, 174
  - service transition 116
- Platform-as-a-Service 23
- platforms 48, 51, 62, 98, 170, 175, 178
- points 82, 96, 106, 169, 171, 178-9
- policies 5, 107, 113, 116, 126, 131-2
  - Service Management 112
- pools, massive 58, 67, 69
- portability 6, 13, 49, 51, 167-8, 174
- ports 93, 96
  - single 96
- PowerPC-based virtualization application 50
- practices, service operation 138
- precautions 74, 77, 90, 93
- prices 80-1
- prioritization 136-7
- processes 8, 10-11, 15, 19, 23-5, 40, 61, 72-3, 80, 105-7, 115-16, 122-5, 131-2, 140-1, 144, 159-60
  - [11]
  - conversion 78, 82
  - time-consuming 63
- processing cores 102
- processing overheads 85
- processing power 8, 12, 14, 22, 41, 68, 73
- processor emulator 51
- processors 22, 24, 26, 28, 31-3, 36, 38-41, 44, 51, 63, 86, 102
  - enabled 36
  - physical 35-6, 38, 40-1, 44, 61
  - single 40, 86, 102
  - virtual 35-6
- processors virtualization tools 41
- processors work 39
- professionals 8, 28, 70, 78, 88, 90, 95, 101, 104-5, 177, 181
- programs 13, 27, 46, 57, 77, 97, 129, 178-9, 181
  - application-level virtualization 56
- project 6, 132, 144, 157, 167-74
- project management 6, 134, 140, 144, 167-73
- project management solutions 167
- project management team 172, 174
- project management tool 171-2
- project managers 6, 169-72, 174
- project plan 170-4
- project team 170, 172-4
- projected service outage (PSO) 134, 139
- protocols 59
- providers 18, 20
- PSO (projected service outage) 134, 139

## Q

QEMU 51  
 QEMU virtualization application 57  
 QoS, *see* quality of service  
 quality 117, 122, 125, 134, 150, 167, 172, 181  
 quality of service (QoS) 125, 146, 154, 159, 162  
 quality service management 129, 134

## R

racks 43, 58, 67  
 RAID 8, 24, 69, 162  
 Rapid deployment of applications 63  
 Really Simple Syndication, *see* RSS  
 reason 11, 40-1, 57, 83, 98, 103, 137, 139, 145, 165-6  
 reboot 66  
 Red Hat 32  
 Red Hat Enterprise Linux 32-3, 55  
 Red Hat Virtualization 4, 32  
 reduced downtime 9-10  
 Reduced operational costs 152-3  
 redundancies 164  
 reinstall 23-4  
 rejection 136-7, 139  
 relationships 6, 139, 144, 168-71, 173  
 release 5, 107-10, 112-20, 122-4, 134, 145, 147-8  
 Release & Deployment 131  
 Release & Deployment Management 107  
 RELEASE & DEPLOYMENT MANAGEMENT 5  
 Release and Deployment 107-8, 111, 119, 140  
 Release and Deployment Activities 115  
 Release and Deployment management 140  
 Release and Deployment Management 139  
 Release and Deployment plans 108  
 Release and Deployment process 112, 115  
 release packages 108, 110, 115, 119  
 release packaging 119  
 Release Policy 5, 110, 113, 115  
 release teams 114, 118  
 release units 109-10, 112  
 reliability 61, 142, 167  
 Request for Changes, *see* RFCs  
 reserves 16, 160  
 resistance 135, 171  
 resource costs 9-10, 20  
 resource requirements 116, 137, 139  
 resource utilization 55, 92-3  
 resources 8, 10-11, 13-14, 18-20, 40-1, 45, 57-8, 62-3, 103-4, 109-10, 113-14, 116, 138-9, 146, 150, 169-71 [16]  
   business s 10  
   consolidate server 88  
   dedicated 16, 160  
   hardware-based 31  
   networked 58  
   organization s 41  
   physical 17, 21, 25, 39, 93, 95, 105  
   physical machine s 13  
   server s 9, 68  
   software-based 46  
   system 39, 41  
 responsibilities 6, 112, 115-16, 122, 131-2, 139-41, 156, 175  
 review 116, 124, 136-7, 140-2, 148, 169  
   formal 124  
 RFCs (Request for Changes) 133-4, 136-7, 141-5, 148

risks 5, 116, 124, 130-1, 133, 137-40, 152, 157, 165, 167-73  
    potential 16, 142, 157, 160  
Road to Storage Virtualization 163  
roles 6, 28, 91-2, 103, 105, 108, 116, 138, 140-1, 149  
routers 68, 96  
RSS (Really Simple Syndication) 170, 173  
RSS technology 170

## **S**

SaaS 18, 20, 24, 178, 183  
SaaS applications 18  
SAN (storage area networks) 43, 58-60, 67-8, 70, 77, 101, 159, 162-4, 166  
SAN appliances 162-3  
SAN technology 68  
scenarios 31, 94, 97-8, 100, 113  
schedule 115, 169-70  
scope 5, 107-8, 117, 128, 130-1, 138, 143, 170, 175  
SCSI 68  
sectors 15, 160  
security 10, 32-3, 38, 42, 70, 72, 79, 91, 117-18, 138, 175-6  
security flaws 70-1, 74  
security measures 70, 74  
Security Monitoring and Virtualization 70  
Server Edition 4, 26, 30  
server hardening 73  
server hardware 38, 82  
server hardware resources 9  
server level 69, 75  
server microprocessors 61  
Server/Workstation virtualization 35  
serverguide/C/virtualization.html 184  
servers 9-11, 16, 20, 31-2, 38-40, 42-3, 67-9, 71, 76-8, 81, 93, 102, 104, 158-60, 162, 166 [8]  
    based 47  
    centralized 38, 64  
    existing 67, 77, 80-1  
    multiple 16, 93, 160  
    single 16, 65, 91, 158, 160  
    virtualized 53, 91, 166  
server s 69  
Service Asset 130  
Service Asset & Configuration Management 131  
Service Asset and Configuration Management 110  
service availability 175  
service availability targets 139  
service catalogue 156  
service change 112, 122, 130, 141  
    authorized 108  
service Change Management process 132  
service change process 130  
service component 114  
Service Console 30  
Service Continuity Management 147  
Service Delivery 177  
Service Delivery Managers 179  
service deployment 114  
Service Design 115  
Service Design documentation 118  
Service Design Package 116  
Service Desk 149-53, 155-7  
SERVICE DESK 6, 15, 176  
Service Desk, centralized 153  
SERVICE DESK, centralized 152, 176

- Service Desk, organization s 157
- SERVICE DESK, virtual 4, 15, 153, 176
- service desk analysts 15
- service desk function 150
- Service Desk Managers 155, 157
- Service Desk Organizational Structures 6
- service desk staffing 6, 155
- Service Desk Types 6, 155
- service desks 155, 176
  - dispersed 154
  - multiple 151
- service disruption 139
  - correct 130
- service group 138
- Service Knowledge Management System 5, 127
- SERVICE KNOWLEDGE MANAGEMENT SYSTEM, *see* SKMS
- Service Level Agreement, revised 139
- Service Level Management and Information Security Management 175
- service levels 114, 151
- service lifecycle 115, 125, 128, 134
- service lifecycle processes 134
- service management 128, 157, 176
  - Web 178
- service management activities 107
- service management implementations 107
- service management perspectives 176
- service management processes 5, 19, 105, 131-2, 175, 181
- Service Managers 142, 157, 179
- service models 138
- Service Operation 108, 122, 124
  - normal 128
- service operations 108
- service operation , normal 149
- service organization 176
- service performance 175
- Service Portfolios 131
- service provider 117, 125, 127, 129, 131, 181
- service provider interfaces 123
- service provider staff 125
- service provision 149
- service quality 125, 128, 130, 143, 146, 157, 162
- service readiness 120
- service rehearsals 119
- service request 133
- service retirement 122
- Service Strategy and Service Design 112
- service teams 109
- service test 116
- Service Transition 124
  - close 116, 124
- Service Transition processes 108
- Service Validation 116, 119
- services 12-14, 18, 21-4, 71, 105-9, 112-13, 119-20, 122-3, 125, 127-30, 138-9, 149-50, 158-9, 175, 177-8, 182-4 [12]
  - 24-hour 176
  - 24-hour follow-the-sun 154
  - changed 109, 113-14, 118, 141
  - client-server-based 53
  - consistency of 113, 153
  - deploy 122
  - desktop virtualization 52
  - live 109, 139

- modified 112, 122
- multiple 138
- production 108
- retired 131
- set 106
- sharing 67
- single 13
- subscription 168
- subscription web 169
- system 33
- technology-enabled 18
- terminal 38
- transfer 122
- Services Staff 141
- services 17
- service 18
- simulation 8
- single physical machine 8, 12-13, 30, 37, 40, 43, 56, 70, 86-9, 94-7, 102, 104, 129
- single physical machine hosting 71
- skill levels 6, 155
- skills 59, 104, 155, 157
- SKMS (SERVICE KNOWLEDGE MANAGEMENT SYSTEM) 5, 127-8
- snapshots 28, 54, 57, 74, 78, 82, 89, 162
- Softgrid Application 4, 33
- Softgrid Client 34
- Softgrid management console 33-4
- Softgrid Management Console 34
- Softgrid package 34
- Softgrid sequencer 33-4
- software 4, 18, 22, 31-2, 39-41, 44-6, 51-2, 61-3, 74-5, 92, 97-9, 101-2, 110-12, 114, 128-30, 174-5 [13]
  - anti-virus 72, 92
- software applications 33, 57, 64, 77, 97, 102, 128-9
- software applications won't work 77
- Software-as-a-Service 24
- software-based virtualization technologies 32
- software conflicts 27, 46, 58, 63-4
- software costs 81, 87
- software environments 84
- software platform 8, 51
- software professionals 92
- software programs 35, 54, 56
- software releases 114
  - definitive 111
- software resources 45
- software support 98-9
- software systems 21-2
- software vendors 89, 98-9, 102
- software vendor's 98, 129
- software virtualization 4, 29, 35
- software virtualization programs 29, 33
- Software Virtualization Solution, *see* SVS
- software virtualization vendor 103
- Solaris 50, 52-3, 55
- solutions 12-13, 18, 34, 37-8, 41, 44-5, 48, 58-9, 63, 66, 75, 84-8, 162-4, 167, 169, 172 [8]
  - best 86, 88-9
  - cloud 167, 174
  - effective 84, 174
  - hosted 81, 84-5, 87, 98, 178
  - robust 34, 41
  - storage-as-a-service 158
- sources 126, 168, 170-1

- multiple 171
- open 6, 26, 32, 48, 81, 94, 167, 171-2
- sourcing, open 172
- space
  - logical 158
  - virtual 27, 29, 33
- spyware 28
- stacks 161-2, 164
- staff 21, 123-5, 128-9, 140-1, 151-3, 156-7, 176
  - service desk 123
- stakeholders 108, 116, 134-5, 140-1, 168-70, 173
- Standard support processes 156
- standardizes 59
- standards 110, 112, 132
- start 78-9, 98, 170-1
- steps 58, 77-83, 90, 132, 136, 161, 163
- storage 17, 39, 42, 59, 67, 70, 77, 81, 107, 109, 158-9, 161-3, 165, 176
  - direct attached 163
  - physical 60, 111, 165-6
- storage area network management tasks 60
- storage area networks, *see* SAN
- storage array 162, 165
- Storage-based virtualization 162
- storage capacities 68, 165, 167
- storage device level virtualization technology 69
  - effective 69
- storage devices 42-3, 67, 69, 101, 159, 165-7
  - multiple physical 67, 69
- storage disks 165
- storage environment 163, 166
- storage infrastructure 60, 159, 161-2
- storage networks 18, 159
- storage resources 18, 20, 58, 68, 165
  - lower level 159
  - networked 58
- storage services 159
- storage solution 159, 161
  - organizations data 59
- storage stacks 162, 164
- storage technologies 59-60
- storage virtualization 4, 6, 15, 32, 42, 59, 67, 103, 160, 163-5
  - incorporated data 60
  - most important feature of 16, 161
- storage virtualization technologies 5, 67, 69, 76
  - adopted data 60
  - deployed 43
  - development of data 60
- storage volumes 42-3, 77, 165, 167
- strain 31, 41, 44, 61, 86, 95
- striping 69, 164-5
- Study Guide applications 42
- Study Guide Bare 84
- Study Guide Blade 81
- Study Guide Block virtualization 17
- Study Guide Container 46
- Study Guide data 18
- Study Guide Definitive Media Library 111
- Study Guide Disaster Recovery 89
- Study Guide effects 63
- Study Guide Environments 117
- Study Guide Essentially 160
- Study Guide Hardware performance 86

Study Guide Hercules Emulator 48  
 Study Guide Increased Security 13  
 Study Guide Information Security Management 107  
 Study Guide Keep 72  
 Study Guide Knowledge 127  
 Study Guide Knowledge Management 125  
 Study Guide Mac-on-Mac 49  
 Study Guide Monitor 74  
 Study Guide Naming Conventions 78  
 Study Guide New 100  
 Study Guide Oracle 50  
 Study Guide Oracle s 75  
 Study Guide Organizations 40  
 Study Guide Overview of Important Steps 116  
 Study Guide Popular application virtualization software 47  
 Study Guide Push 114  
 Study Guide Reduced Management and Resource Costs 10  
 Study Guide Release Package 110  
 Study Guide role of technical support staff 129  
 Study Guide Softgrid Management Console 34  
 Study Guide Sun 52  
 Study Guide system 90  
 Study Guide True virtualization 41  
 Study Guide Using 171  
 Study Guide Virtual Desktop Infrastructure 38  
 Study Guide Virtual Private Servers 53  
 Study Guide virtualization application 85  
 Study Guide virtualization solutions 82  
 Study Guide Win4Lin 57  
 Study Guide work 66  
 subscription 18-19  
 Sun Microsystems 52  
 Sun Microsystems 52  
 Sun xVM Server 52  
 Sun xVM VirtualBox 52  
 Super Users 156  
 Super Users 156  
 suppliers 21, 108-9, 124, 131, 134, 141, 144, 169  
 support 21-2, 28, 32, 36, 43, 47-8, 50, 52-6, 58, 61, 68, 74-5, 96-9, 129-30, 149-50, 153-7 [7]  
     hardware-assisted virtualization 63  
     multiple host 164  
     quality of 157  
     virtualization software platform 86  
 support Change Management 131  
 support issues 5, 97, 117  
 support post-deployment 122  
 support servers 32  
 support staff 108, 120  
 support virtualization 77  
 support virtualization technology 97  
 SVC 28  
 SVISTA 51  
 SVS (Software Virtualization Solution) 27, 29, 33, 47  
 switches 13, 29, 68, 88, 91, 96, 106, 164  
 Symantec s Altiris Software Virtualization Solution 33

## T

tape 15-16, 160  
 tape media 16, 160  
 tape media virtualization 16, 160  
 TCO (Total Cost of Ownership) 112  
 teams 103, 122, 149, 171-2

- technical support 97-8, 128-9
- technical support staff 97-8, 128-9
- techniques 16, 42, 69, 128, 161, 164-5, 167, 171, 173
- technologies 9, 31, 33, 35, 38, 46, 51-2, 58, 60, 66, 68, 83, 97-9, 141, 149-50, 153 [6]
- tedious 82
- templates 90, 173
  - virtual 90
- test 13, 29, 31, 77, 79-80, 88, 96, 98-9, 108-9, 114, 116, 118-19, 134, 139, 144, 166
- test environments 37, 119, 138
- theartofservice.net/index.php 17
- theme 172
- thin client technology 53, 57, 64-5
- thin-client technology 93
- ThinApp 54, 62-3
- ThinApp applications 62
- Thinstall virtualization 62
- Thinstall Virtualization Suite 5, 62
- time zones , multiple 153-4
- tools 41, 51, 53, 76, 80, 91, 104, 119, 127, 171-4
  - service desk 15
  - virtualization enhancement 41
- Total Cost of Ownership (TCO) 112
- track 91-2, 107
- training 104-5, 115, 117-18, 156
- transfer 116, 122-3
- transition 80, 105, 109, 122, 124, 134
- transition processes, surrounding Service 112
- troubleshooting 97, 128-9

## U

- Ubuntu 4, 26, 32, 55
- understanding 125-6, 129-30, 160, 171
  - Business 156
- Updated service catalogue 112
- updates 34, 39, 52, 90, 114, 127, 172, 174
- upgrades 32, 34, 95
- uptime requirements 95
- usage 19, 174
- USB ports 79, 93
- user base 113
- user environments 36, 38
- user groups 15
- user interface, graphical 26, 34, 47, 51-2, 57
- users 18, 29-30, 37-8, 41-2, 46-7, 52-4, 56-7, 61-2, 64-5, 67-8, 108-9, 114, 122-4, 149-53, 155-6, 167-75 [14]
  - associated 110
  - expected 114
  - home 37, 41, 61
  - multiple 68, 168
  - End 112
  - Mobilizing 120
- users processes 41
- user s 33, 54, 93
- Using virtualization management tools 72
- Using Xenocode 57

## V

- validation 118-19
- variances 16, 106, 143, 160
- VDI (Virtual Desktop Infrastructure) 37-8, 52
- vdisk 158
- VDM (Virtual DOS Machine) 54



- vendors 59-60, 78, 80, 83, 85, 89, 97-8, 128-9
  - application s 77
- veredition/technologies/virtualization 184
- versions, latest 28, 32, 64
- VirtIO 27
- Virtual Desktop Infrastructure, *see* VDI
- Virtual DOS Machine (VDM) 54
- virtual environment 44-6, 55, 62, 74, 77, 80, 83-4, 86, 93, 102, 166
  - isolated 24-5
- virtual hosting environments 45
- Virtual Iron 54
- Virtual Machine, based 66
- virtual machine clones 13
- virtual machine data 13, 93
- Virtual Machine Disk Format 55
- virtual machine files 37, 82
- virtual machine log files 74
- Virtual machine management software 10
- virtual machine management solution 87
- Virtual Machine Managers 39, 55
- Virtual Machine Monitor (VMM) 24, 39
- virtual machine software 23
- virtual machine templates 90
- virtual machines 10, 12-14, 20-5, 28-30, 32-3, 36-8, 43, 50-1, 53-7, 71-4, 76-80, 84-7, 89-97, 99-101, 103-4, 107 [5]
- virtual machines files 89
- virtual machines OS 24
- virtual machines role 91
- virtual machines 11
- virtual private servers 47, 52
- virtual servers 9, 31, 48, 77, 85-7, 93-4, 101
  - multiple 28, 40, 71, 88
- virtual Service Desk structure 153
- VirtualBox 55
- virtualization 4-6, 8-9, 11-15, 19-26, 31-3, 35-42, 44-6, 59-60, 64-6, 69-73, 83-5, 88-90, 95-100, 103-7, 128-9, 158-67 [6]
  - bare metal 46, 83
  - database 16, 161
  - desktop 4, 37, 75
  - desktop-based 104
  - efficient 39
  - file system 16, 161
  - file system level 70
  - hosted 46, 83-4
  - in-band 162-3
  - network-based 162, 164-5
  - operating-system level 36
  - out-of-band 162-3
  - partial 23
  - processor 4, 8, 40, 75
  - resource 4, 39-40
  - software-based 32, 38, 75
  - Files systems 161
- Virtualization and Storage Management 158
- virtualization application 12, 24, 31, 46, 48-9, 51, 55, 74, 81, 84, 159
  - desktop 52
  - free server 30
  - hosted 46-7, 49, 53-5
  - longest running modern day 30
  - popular 30
  - popular hosted 46
  - portable 51

- virtualization approaches 164
- virtualization-based strategies 35
- virtualization companies 30-1
- virtualization density 4, 43
  - term 43
- virtualization engines 163-4
  - multiple 164
- virtualization environments 129
- virtualization form 16, 161
- Virtualization Happen 161-2
- virtualization implementation 163
- virtualization layer 163
- Virtualization levels 5, 69, 75
- virtualization management application 47, 49
- virtualization management software 91
- virtualization management tools 83, 89, 91
- virtualization manager 54
- virtualization market area 28
- virtualization processes 4, 19, 26, 39, 82, 94, 103, 163
- virtualization programs 30
- virtualization project 21
- virtualization security policies 70
- virtualization software 4, 8-9, 16, 27, 35, 39, 42, 46, 48, 50, 61, 80-6, 97, 99, 103-4, 158 [3]
  - bare metal 52, 82
  - deploy 98
  - efficient 85
  - fundamental 98
  - hosted 52
  - installing 37, 84
  - popular 47
  - popular bare metal 46
- virtualization software solutions 60
- virtualization software vendors 82, 99, 103
- virtualization solutions 31, 33, 36, 40, 53, 56, 66, 81-6, 88-9, 93-5, 99
  - application-based 26
  - data server 60
  - data storage 58-60, 69
  - desktop 93
  - efficient 41, 85
  - hardware-assisted 60, 100
  - hybrid 36
  - money-saving 95
  - monitor 74
  - operating system-level 55
  - operating-system-level 56
  - popular bare metal 82
  - resource-intensive 37
  - software-based 39, 88
  - tested 28
  - third 84
  - traditional 38
- virtualization solutions purpose 88
- Virtualization Specialist 182
- Virtualization Specialist Exam Tips 7, 182
- Virtualization Specialist Level Complete Certification Kit 4-185
- virtualization stability 39
- virtualization strategies 47
- virtualization system 158, 164
- virtualization techniques 8
  - common 162
  - network-based 165
- virtualization technologies 4, 8, 11, 26-31, 35, 38-9, 41, 43, 46, 48, 65, 70-1, 76, 83, 96-9, 128-9 [2]

- strong 28
- Virtualization Technologies Review Questions 5, 75
- virtualization technology built-in 26
- virtualization technology matures 74
- virtualization tools 78, 87
- virtualization type technologies 35
- virtualization vendors 97
- virtualized areas 27, 32
- virtualized blade server environment 81
- virtualized environment 5, 13-14, 39, 42, 44, 55, 70-1, 73, 76-7, 79, 83-7, 89-92, 94-5, 97-103, 128-9, 163-5 [1]
  - new 80, 96, 103-5
- virtualized guests 32
- virtualized infrastructure 10, 70-1, 82-3, 86, 91-2, 95-6, 100-2, 104
- virtualized server environment 28
  - standard 81
- virtualized software environment 26
- virtualized solution 63, 69, 83
- virtualized storage area 40, 42-3
- virtualized technology 4, 12, 26, 43
- virtualizing 158-9
- Virtuozzo 55-6
- virus 10, 13, 72-3, 91-2
- virus infection 72
- Vista 55, 89
- VMDK 55
- VMM (Virtual Machine Monitor) 24, 39
- VMware 8, 12, 26, 28-31, 53-5, 60, 62, 83
- VMware applications 30
- VMware ESX Server 30-1, 53, 82
- VMware ESX Server application 30
- VMware ESXi 30
- VMware Fusion 30, 53
- VMware Player 30, 53
- VMware Server 30-1, 53
- VMware ThinApp 54, 62, 64
- VMware ThinApp applications 62
- VMware Workstation 30-1, 54, 60
- VMware's Converter Tool 78
- VMware's ThinApp application 62
- VMware's Workstation solutions 52
- volumes 16, 152, 156, 160, 167-8
- Vx32 56
- Vx32 application 56

## **W**

- waste 10-11, 41, 77
- web 167-9, 174
- web applications 167-8, 171, 173
- web hosting environments 46
- web page 169-70
- what's 139
- Win4Lin 56
- Windows 11-12, 28, 30, 45, 48-55, 57, 64, 66, 81, 99
- Windows applications 42, 54, 62
  - normal 57
- Windows-based operating systems 55-6
- Windows Server 4, 28, 48, 81, 99
- Windows Server Virtualization 28, 48
- Windows Virtualization 5
- Windows Virtualization on Fedora 66
- Windows Vista 57, 64, 89

Windows XP 55-6, 64  
wisdom 126-7  
won t 87-8  
work 11, 26, 32, 39, 41, 44, 57, 60-1, 63-5, 76-9, 96, 99-100, 103-4, 163-4, 169, 171-2 [5]  
work orders 136, 140  
workers 103-4  
workload 45, 60, 84, 86, 106, 119, 164-5  
    based virtualization 31  
    processor s 93  
workstations 10-11, 34, 113  
workstations/servers 8  
writing 16, 160  
WSVs 28  
www.artofservice.com.au 181, 185  
www.redhat.com/rhel/virtualization 184  
www.theartofservice.com 185  
www.theartofservice.org 185  
www.vmware.com/virtualization 184

**X**

Xenocode 57-8  
Xenocode applications 58