Multitenant Utility Computing on IBM Power Systems Running AIX

- Implement best practices for running IBM AIX in service provider environments
- Leverage the latest IBM PowerVM virtualization technologies
- Use management and monitoring tools provided with IBM AIX Enterprise Edition

Chris Almond
Paulo Cesar Chiquito
Carlos Henrique Fachim
Sehee Kim
Masahiko Okajima
Pauli Rämö

ibm.com/redbooks

Redbooks

International Technical Support Organization

Multitenant Utility Computing on IBM Power Systems Running AIX

February 2009

SG24-7681-00

Note: Before using this information and the product it supports, read the information in "Notices" on page vii.

First Edition (February 2009)

This edition applies to IBM AIX 5.3, IBM AIX Version 6.1 running on IBM Power Systems, and to HMC Version 7, Release 3, Modification 30, Service Pack 2.

Contents

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AFS®
AIX 5L™
AIX®
BladeCenter®
CICS®
DB2 Universal Database™
DB2®
eServer™
Focal Point™
GPFS™
HACMP™
IBM Systems Director Active Energy Manager™
IBM®
Micro-Partitioning™
NetView®
POWER Hypervisor™
Power Systems™
POWER4™
POWER5™
POWER5+™
POWER6™
PowerVM™
POWER®
pSeries®
Redbooks®
Redbooks (logo) ®
System i®
System p5®
System p®
System x®
System z®
Tivoli®
TotalStorage®
WebSphere®
Workload Partitions Manager™

The following terms are trademarks of other companies:

Acrobat, Adobe, and Portable Document Format (PDF) are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

Oracle, JD Edwards, PeopleSoft, Siebel, and TopLink are registered trademarks of Oracle Corporation and/or its affiliates.

SAP R/3, SAP, and SAP logos are trademarks or registered trademarks of SAP AG in Germany and in several other countries.

VMware, the VMware "boxes" logo and design are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions.

Java, JDBC, Solaris, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Excel, Microsoft, SQL Server, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Preface

This IBM® Redbooks® publication presents concepts, considerations, and high level deployment examples for technical professionals who design and operate multitenant utility computing environments hosted on IBM System p® and IBM AIX® Enterprise Edition. AIX Enterprise Edition brings together IBM enterprise management capabilities from Power Systems™, IBM AIX, and Tivoli® Software to provide a powerful set of integrated functions for infrastructure management in multitenant utility computing environments.

In this book, we focus on the following topics:

- Using System p, IBM AIX 6, and PowerVM™ virtualization technologies
- Provisioning in a multitenant computing environment on AIX and Power Systems
- Monitoring resources in a multitenant environment
- Accounting and chargeback in a multitenant AIX environment

AIX Enterprise Edition brings together several IBM products: AIX 6 operating system, IBM PowerVM Workload Partitions Manager™, Tivoli Application Dependency Discover Manager, Tivoli Monitoring, and IBM Usage and Accounting Manager Virtualization Edition for Power Systems. In addition to these products we also provide a system allocation example using Tivoli Provisioning Manager.

The team that wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), Austin Center.

Chris Almond is an ITSO Project Leader and IT Architect based at the ITSO Center in Austin, Texas. In his current role, he specializes in managing technical content development projects focused, Linux®, System p/AIX systems engineering, various IBM Software Group products, and innovation program topics. He has a total of 18 years IT industry experience, including the last eight with IBM.

Paulo Cesar Chiquito is a Security Specialist in the Server and Technology Group in IBM RTP, USA. He has 15 years of experience in AIX, Solaris™, Linux, Samba, DCE and AFS®, Perl and Siebel®. His areas of expertise include IT Security, Distributed Filesystems, System Administration, and automation using the Perl language. He is a certified Siebel Developer and a Certified Advanced Technical Expert Specialist.

Carlos Henrique Fachim is a Advisory IT Specialist working in the Integrated Technology Services in IBM Brazil. He joined IBM 10 years ago, working as an AIX Support Specialist. His expertise areas include AIX (also DUMP and Performance Analysis), HACMP™, Power Systems and PowerVM features, SAN and TotalStorage® products (DS6000/DS8000, SVC and TPC). He also has experience designing and implementing highly virtualized and highly available environments. He holds a Master’s degree of Computer Science from Mackenzie University (Sao Paulo, Brazil). He is a Certified Advanced Technical Expert Specialist.

Sehee Kim is a Senior IT Specialist in Technical Sales, IBM Korea. She has eight years of experience in AIX, Power Systems, HACMP, GPFS™ and two years of experience in application software development before joined IBM. She holds a Master’s degree in Computer Science from Seoul National University, Korea. Her areas of expertise include benchmark testing, C/C++ programming, porting software, deep computing (HPC), and UNIX® systems. She is also an AIX Certified Advanced Technical Expert.

Masahiko Okajima is a Advisory IT Specialist working for the IBM Global Services Japan Solutions and Services company (ISOL). He has been with IBM for 10 years, and has experience in AIX and IBM Power Systems, storages, networks, and various middlewares such as WebSphere® Application Server, DB2®, Tivoli products, and more. He is also an expert in designing and implementing highly available systems, backup and recovery systems, and in monitoring systems.

Pauli Rämö is a Senior IT Specialist working for the Integrated Technology Services in IBM Finland. He has 16 years of experience in AIX on Power Systems and HACMP. He has eight years of experience in enterprise storage solutions, including disk, tape, and SAN technologies. He also has six years of experience in designing and implementing multitenant solutions based on Power Systems LPAR technology. He is an IBM Certified IT Specialist, and has several product certifications on AIX, SAP® R/3, and ITIL®. He is the co-author of four SAN IBM Redbooks publications.

Acknowledgements

This book would not have been possible to write without the generous support of many IBM employees around the world. For their contributions, guidance, patience, and technical feedback in support of this project, we gratefully acknowledge the following IBM employees:

Sandy Amin
Mark Ball
Bruno Blanchard
George Bonifant
Ed Geraghty
David Hepkin
Greg Howard
Bob Kovaks
Shawn Mullen
Edward Shvartsman
Mary Wigham

The team would also like to acknowledge the support for this project provided by:

Jay Kruemcke, IBM AIX Offering Manager
Octavian Lascu, ITSO Project Leader
Scott Vetter, ITSO Project Leader

Become a published author

Join us for a two- to six-week residency program! Help write a book dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You will have the opportunity to team with IBM technical professionals, Business Partners, and Clients.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you will develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  ibm.com/redbooks

- Send your comments in an e-mail to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

1

Introduction

This chapter provides an introduction to utility computing, multitenancy, how it evolved and why it is advantageous to your business. We also define the concepts related to multitenant utility computing and factors to consider when you choose the level of resource sharing of your architecture.

This chapter contains the following topics:

1.1 Reasons for utility computing

The first questions we try to answer are:

- Why move the IT delivery to a model similar to utility computing?
- What are the advantages of the utility model?

The utility computing model can deliver:

- The right IT solution or IT service
- Terms that are acceptable to the user, as defined by the service level agreement (SLA)
- The most effective use of resources, such as by increasing efficiency
- The flexibility to adapt to business changes quickly
- A cost effective approach

To better understand how utility computing fulfills the needs and answers our questions, we have to understand what the consumers of IT services want.

We can divide IT delivery in three parts:

- Business

  The business is mainly concerned about cost, how well the IT solution fits the business model and how quick can IT change to accommodate business changes.

- IT organization

  The IT organization is the one that delivers the IT services.

- IT consumer or user

  The user does not care about how the IT service is implemented, as long as it can have an reliable service that fulfills its needs. The user only pays attention to how IT is services are performed when service breaks.

The IT organization is under pressure from both sides. The organization is required to lower costs of the business while at the same time the user wants increased performance and functionality.

Utility computing is as much of a process and principle as it is a technology. Utility computing encompasses the following concepts:

- Deliver individuals services
- Provide the right service at the tight time in the right amount
- Increase accountability for the service cost, delivery and performance

- Optimize performance, reliability and cost
- Respond quickly to changing business requirement

An IT organization has to cycle though the following steps to successfully navigate utility computing:

1. Constantly analyze and find the real business needs.

   A crucial task is to stop and reassess whether IT is aligned with the business. When was the last time IT stopped and really assessed its role in business?
2. Consolidate and virtualize.

   Consolidation and virtualization are two technologies to increase standards of the IT delivery. The goal is to reduce costs and make automation and standardization easier.
3. Adhere to standards.

   To be able to realize economy of scale, following standards is critical. Custom solutions should be avoided, when possible.
4. Deliver everything as a service.

   Delivering everything as a service is simply making sure the way you interface with your users follows a standard.
5. Automate all operations.

   Automation is the second driver for lowering costs.

Electrical power is a utility supplied by the electrical company, utilizing standardized transport, delivery and metering for the utility (the power). Similarly, utility computing is a way to provide computing power (for example CPU cycles and adjacent services, storage and access to data) by standardizing service delivery, metering, and billing.

1.2 The history of utility computing

Although utility computing is a new term, the concept has a long history. Early client server environments used mainframes that were big, reliable, and required highly trained operators. They were shared across a large user base to justify the investment. The mainframe users presented job cards. The jobs were processed over a period of time and the users would receive the results, usually on a printed paper.

Billing was based on resource usage; users were charged for their share of mainframe usage.

That mode of operation, we can argue, represents the utility computing mode of delivery, as follows:

- Standard interface exists.
- Billing for usage is implemented.
- User provides input, and receives output in standard formats.
- Dynamic capacity is provided.

So, computing started with a model very similar to utility computing, maybe not as flexible, but common on many features.

With the advent of the minicomputer, and later the personal computer, and the commoditization of computer hardware, IT moved away from the early version of the utility computing delivery model.

With the inexpensive hardware, everyone had one or more computers with lots of custom software solutions. As businesses started adopting the same solution, IT departments now had to support a myriad of computer platforms, different software versions, and different applications. The cost of managing that infrastructure increased exponentially as the complexity and variety increased.

Some elements of utility computing are still present, such as inexpensive, expandable, standard, and powerful hardware. However, the missing link that could have enabled utility computing was a fast, cheap, and easy way to connect and use the services provided by that hardware.

The 1980s and 1990s brought to light new technologies that enabled the utility computing delivery: the Internet and the World Wide Web. The Internet provided a new way to interconnect machines across the globe. The World Wide Web provided a set of technologies to exchange information quickly and in a standard way. The next missing link was a set of applications that could make use of the utility infrastructure.

Today, we have all the elements to successfully enable utility computing:

- Communication standards define common ways to send and receive information.
- Flexible, easy to scale, cost effective hardware.
- Users pay only for what they use.
- Resources are centralized and managed by the service provider.
- Users do not have to be concerned about how services are implemented.

More recently, certain concepts, if not the same, are very similar to utility computing. They try to be more specific about areas of utility computing, such as:

- Software as a service (SaaS)

  The software application is delivered as a service. Instead of hosting an application the user can buy application time.

- Platform as a service (PaaS)

  The platform is delivered as service. Instead of installing and maintaining the OS and middleware, the user can buy the platform as a service and install the application on the service.

Utility computing is again a reality, enabled by a different set of conditions and being used by an increasing number of companies.

1.3 Multitenancy, virtualization and utility computing

One of the main driving factors for utility computing is the cost savings that can be realized. The savings are mostly based on:

- Economies of scale

  Duplicating the same solution across a wide range of tenants can mean a quicker Return on Investment (ROI), reducing overall operational costs for service providers. It is also easier to automate the management of a widely replicated solution.

- Automation

  To reduce complexity and administrative costs automation is used to manage utility computing environment. Automation reduces costs by allowing fewer IT personnel to manage more systems.

- Virtualization

  Virtualization is a broadly used term that refers generally to the abstraction of computational resources. Virtualization can occur a the platform, specific resource, or application level. Virtualization enables fast response to change, allows savings by increasing resource utilization and makes the utility computing service provider more flexible to changes.

- Multitenancy

  Multitenancy refers to the sharing of infrastructure between several users. From the service provider point of view It can optimize resource utilization and reduce capital cost.

Not all situations call for virtualization or multitenancy. In some cases potential tenant requirements may require dedicating system resources into a traditional single tenant architecture.

One usual objection to multitenancy is the perception that the security isolation of single tenant architectures is better than multitenant. In Chapter 2, “Multitenancy, PowerVM, and partitions” on page 17 we analyze 4 multitenant cases and their security implications.

In the following sections we will discuss virtualization and the different levels of multitenancy.

1.4 Multitenancy

The term *multitenancy* is a term borrowed from the real estate industry. Traditionally, buildings were used by a single tenant and when buildings started being shared by multiple tenants, the term multitenant was coined. In a multitenant environment tenants would have a private space and common space, shared among all tenants. By sharing resources and creating standard offerings, multitenancy reduces cost and improves efficiency of real estate operations.

1.4.1 Overview of multitenancy

When using the term multitenant, in a very broad sense we mean the ability to provide computing services to multiple customers by using a common infrastructure and code base.

A multitenant architecture also uses virtualization to increase resource utilization, load balancing, scalability, and reliability. Automation is used to reduce complexity, decrease operation costs, and increase provisioning speed.

A typical user (or tenant) of a multitenant system is unaware that the common infrastructure and code base is being shared across multiple customers. The multitenant architecture provides separation of data between customers, enforces security boundaries and provides a mechanism for billing based on resource usage.

A good example of multitenancy is an independent software vendor (ISV) that hosts the databases of two separate companies by actually using the same DB2 database, each one in a separate schema. All resources are shared, security is maintained, usage is recorded but customers are unaware of each other.

With multitenancy, we have to define at what level the multitenancy is applied. Levels can include:

- Application level

  Multiple tenants use an application. The application provides logical separation between users, access controls, and customization. An example is a WebSphere application.

- Middleware level

  Multiple applications use the same middleware. The middleware provides logical separation, access controls, and resources. An example is several instances of WebSphere connected to the same DB2 database Instance. DB2 assures that each application data is separated and protected from each other.

- Operating system (OS) level

  Multiple middleware runs under the same OS. The OS provides access controls, logical separation, and resources to the middleware. An example is one OS running multiple DB2 databases under different instance ids. The OS implements a mechanism to protect data from each DB2 instance.

- Hardware level

  The hardware provides logical separation, access control and resources to each OS instance. Each OS instance is considerate a tenant. An example is VMWare ESX images running on the same machine. VMWare provides logical separation, access controls, and hardware resources to each OS.

Figure 1-1 on page 8 shows a graphical description of the levels where you can have multitenancy and the one that applies to this book.

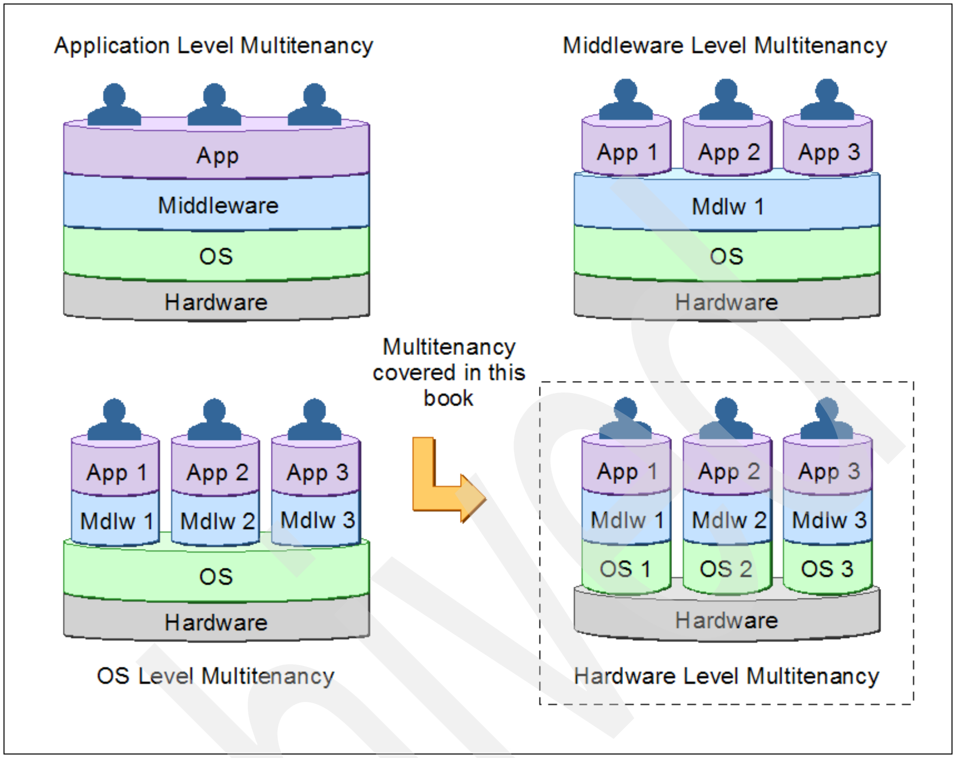

Figure 1-1 Multitenancy at different levels

1.4.2 Multitenancy resource sharing levels

Multitenancy can be realized with different degrees of resource sharing. The most typical components that can be shared across multiple tenants are:

- Disk or tape storage

  Virtualized storage provides several mechanisms to allow the sharing of disk or tape storage across several computer systems.Virtualized storage is usually delivered by devices with virtualization capacity through a storage area network (SAN), but can also be delivered by a direct connection to the host machine.

- CPU processing

  Traditional multiuser systems provide several mechanisms to share CPU processing across several users. Virtualization has added another layer of abstraction, by allowing the creation of virtual instances of the hardware. Each virtual hardware instance can have a virtual CPU that is mapped to a slice of the CPU processing power.

- Network bandwidth

  Network bandwidth is a resource that is traditionally shared. Multitenancy can be implemented with different levels of traffic separation, such as VLANs, IP subnetworks, and network firewalls.

Based of the number of components that can be shared, defining a resource sharing spectrum line, as defined on Figure 1-2, is helpful. The line depicts the amount of resource sharing for multitenant architectures. Architectures on the left have minimal resource sharing, and architectures on the right have a greater amount of resource sharing.

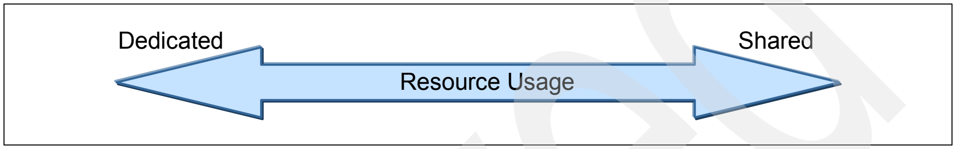

Figure 1-2 Resource sharing spectrum

To increase savings on your capital costs you want to maximize the sharing of your computer resources, however increasing the amount of resource sharing can have several side effects that we explore in more detail in the next section.

1.4.3 Factors that affect multitenancy resource sharing

When you start sharing computer resources across multiple tenants, consider the following factors:

- Cost or Investment

  Increased resource sharing decreases your cost by spreading it across several tenants. Instead of preparing for the worst case workload for a single customer, you may pool all extra capacity together and utilize it to serve individual tenants.

- Complexity

  Multitenant architectures with high levels of resource sharing are intrinsically more complex because they have to provide for customization, additional security layers, monitoring, auditing, metering, and billing.

  The added complexity has to be balanced with a greater use of automation for deployment and management of tenant resources.

  Figure 1-3 on page 10 compares the cost and complexity with the resource usage spectrum.

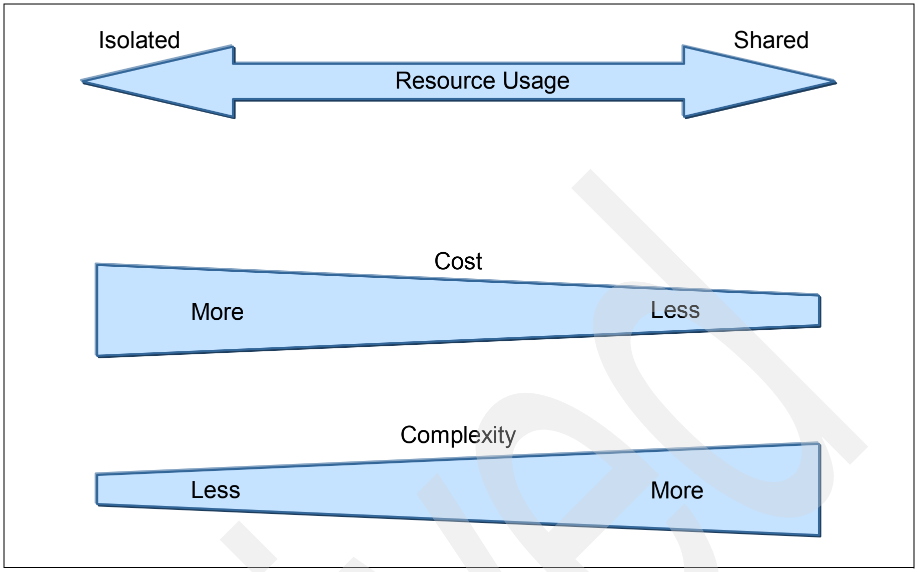

Figure 1-3 Cost and complexity in the resource sharing spectrum

- Security isolation

  Several proven technologies are used to make sure that tenants do not cross security boundaries. However, in cases when the highest security level is required, you might want to physically separate resources that are used by different tenants.

- Performance isolation

  When sharing resources across multiple tenants, make sure that one tenant load does not affect performance of other tenants. The affected resources can be I/O, CPU usage, or network bandwidth. See Chapter 2, “Multitenancy, PowerVM, and partitions” on page 17 for more information.

- Resource utilization

  A multitenant architecture with a greater degree of sharing presents a greater degree of resource utilization. Virtualization functionality can also be used to move additional resources where it is most needed and allows a greater granularity when assigning resources, compared to dedicated resources.

  Isolated architectures have to be sized for worst-case scenarios and can have several pockets of underutilized resources.

  Figure 1-4 on page 11 compares security and resource utilization with the resource sharing spectrum.

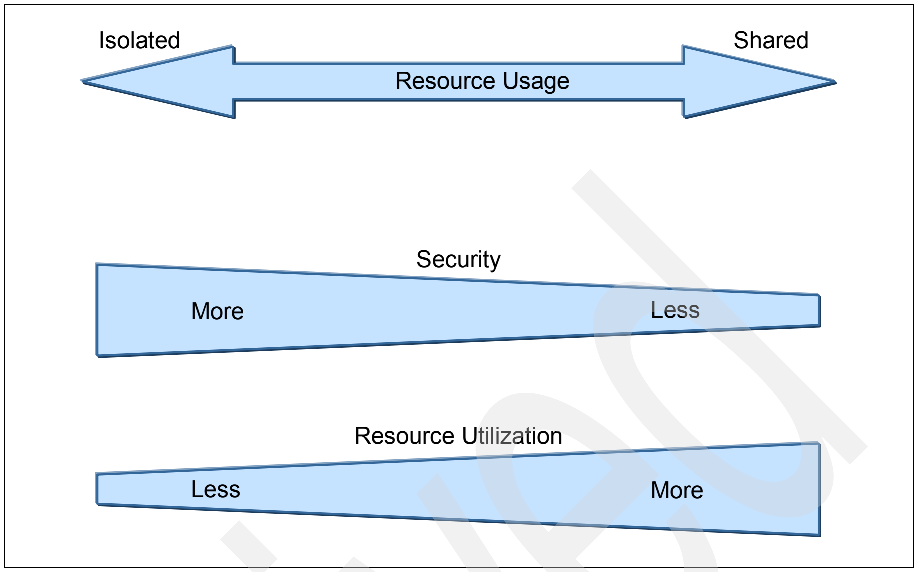

Figure 1-4 Utilization and resource sharing

- Scalability

  Multitenant architectures with high levels of resource sharing are designed to be more scalable. The virtualization capabilities of the underlying hardware can also be used to quickly increase capacity dynamically.

- Flexibility

  Like scalability, flexibility by design must be lot greater in a multitenant architecture with high levels of resource sharing, because of the use of virtualization technologies.

  Figure 1-5 on page 12 compares scalability and flexibility with the hardware sharing spectrum.

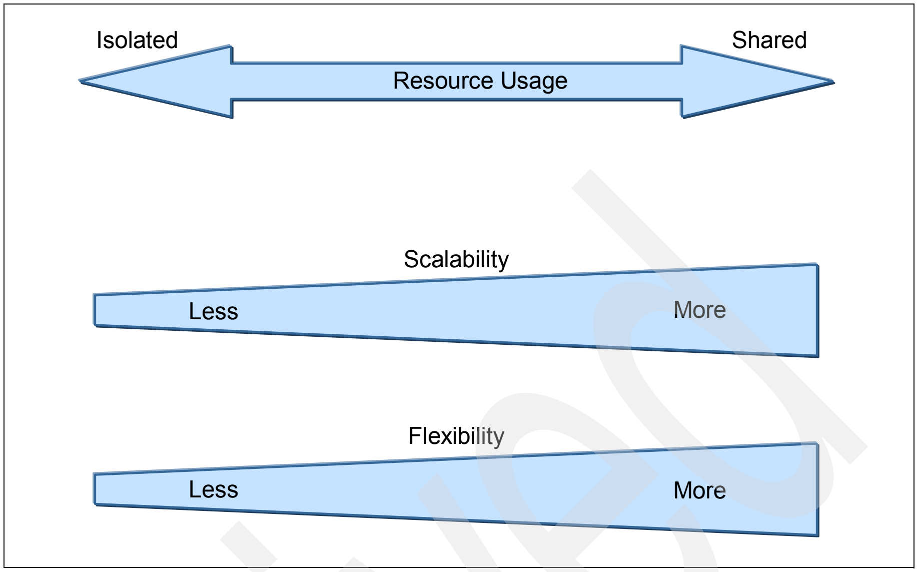

Figure 1-5 Scalability and Flexibility on the resource sharing spectrum

- Availability

  Availability is a system's capability to mask or reduce both planned and unplanned downtime. In a multi-tenant environment, one of the most important availability issues is the capability to provide isolation between tenants. We could argue that in an architecture with shared resources more tenants are affected by a resource downtime. On the other hand, when sharing a resource, you reduce the number of parts subject to failure and you can spend more attention and effort on increasing the reliability of the shared resource.

  For this analysis, we consider that availability is not affected by the level of resource sharing.

- Tenant profile

  The tenant profile is an important factor to consider when choosing the level or resource sharing. Most likely, larger tenants require more dedicated resources, and might want a higher level of customization and added value services. Smaller tenants, on the other hand, might be more concerned about costs.

  Another dimension to consider is the type of tenant workload. To maximize resource utilization, pooling together similar sized tenants with diverse workloads might be better. For example, hosting all tenants performing scientific computations on the same hardware pool can quickly exhaust all the CPU resources.

Grouping similarly sized tenants that combine to create heterogeneous workloads can increase resource sharing and tends to balance resource utilization.

Figure 1-6 relates the resource sharing spectrum line with all the tenant profiles previously discussed.

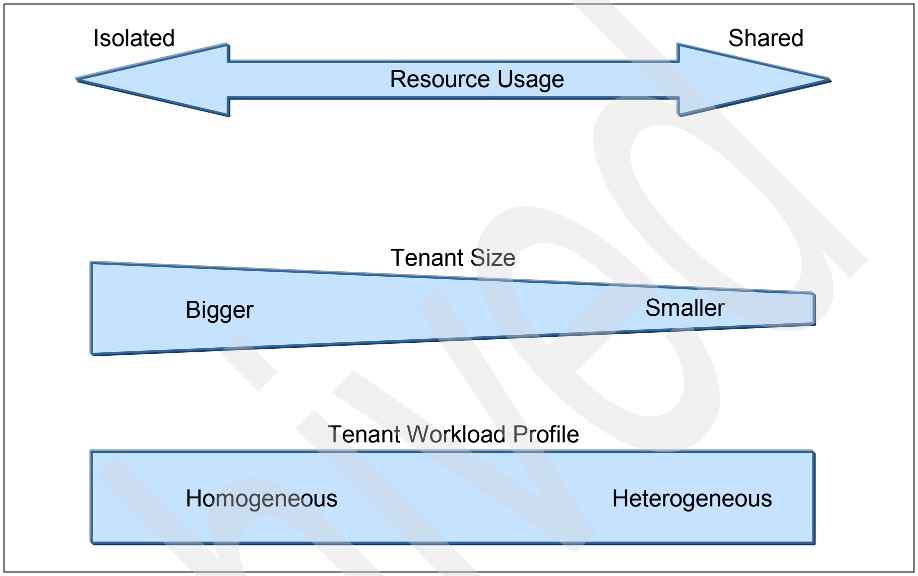

Figure 1-6 Tenant size and workload profile on the sharing spectrum

1.4.4 Multitenancy challenges

When talking about multitenancy, we have to be aware of the challenges that are involved with hosting multiple tenants on the same server. As the number of tenants increases, we face the following challenges:

- Management

  As the number of hardware components increases, management costs also increase. The traditional way of adding capacity for new tenants has been to acquire new hardware components, which increases the diversity of the data center, increases power consumption, and makes management of the hardware even more challenging.

- Provisioning

  Traditionally provisioning has been a very slow process. The process usually involves a capital planning phase, acquisition, delivery, hardware and software installation, and customization. This old model is a big road block for multitenancy, because it slows the growth, and hampers flexibility and capacity growth.

- Complexity

  As the data center grows in size and diversity of equipment, the complexity of keeping the configuration information, its dependency information, and the applications design up-to-date grows.

- Power Usage

  The increased power requirements of new computer components and the exponential use of computing resources has recently escalated power usage on data centers. To make matters worse, increased power consumption is directly related to cooling requirements, which increases the total power usage.

  Power and cooling needs can be another factor in constraining the growth and flexibility of multitenancy.

- Billing or chargeback

  As we move to a per-tenant model, we have to be able to collect sufficient resource metering information so that we are able to provide detailed and accurate billing.

Virtualization is a key technology to address these challenges. In the next section, we explore how virtualization can help you address these challenges.

1.5 Virtualization

Virtualization enables the sharing or the aggregation of physical resources and the creation of logical resources that are used by a guest operating system. The guest OS is not usually aware that it is using virtual resources or how these resources are implemented.

Note: Virtualization in this context can be used in two ways. One way is to divide a resource into smaller pieces, such as a powerful processor in several logical smaller processors. Another way is to combine smaller resources in a larger virtualized one, which is similar to joining several small disks in one large virtual disk. Also, virtualization can provide for uniform access to resources.

1.5.1 Virtualization advantages

Several advantages of virtualization are:

- Improved hardware utilization

  Server hardware resources can be more fully utilized, by being dynamically partitioned.

- Hardware consolidation

  Instead of having multiple small machines with lots of spare capacity, all the workload can be consolidated on a small number of machines, allowing you to experience a rapid return on investment (ROI).

- Increased flexibility

  Virtualization supports the pooling of resources that can be managed centrally through an enterprise hub to better support changing business requirements dynamically.

- Enabled access through shared infrastructure

  Virtualization provides a resilient foundation and shared infrastructure that enables better access to infrastructure and information in support of business applications and service-oriented architectures (SOA).

- Rapid provisioning

  Virtualization can enable rapid infrastructure provisioning, down to the order of minutes, compared to the order of days in a non-virtualized environment.

- Disaster recovery

  Using the rapid provisioning features, additional instances of the application being recovered can be deployed very quickly. An image copy of the instance can also be kept in order to speed up recovery.

- Reduced power and cooling requirements

  By virtualizing your resources, you can consolidate the load on fewer machines, and reduce overall power usage and cooling requirements.

2

Multitenancy, PowerVM, and partitions

In this chapter we introduce the PowerVM virtualization technologies and how they can be used on a multitenant environment. Considerations for using Workload Partitions and partition mobility are also discussed. Finally, we present several tenant scenarios along with considerations you should take in account when designing a multitenant architecture.

This chapter contains the following topics:

2.1 PowerVM Virtualization technologies

> **Note:** This is a brief introduction of the PowerVM virtualization technologies. For details, see *PowerVM Virtualization on IBM system p: Introduction and Configuration Fourth Edition,* SG24-7940, *PowerVM Virtualization on IBM System p: Managing and Monitoring,* SG24-7590, and *Introduction to workload Partition Management in IBM AIX Version 6.1,* SG24-7431.

PowerVM is a brand that includes a range of hardware and software features, which allow a system to be more flexible in adapting various workloads.

PowerVM is offered in three editions:

- PowerVM Express Edition

  Offered exclusively on the IBM Power System Model 550 and IBM Power 520 Express servers, it is designed for users who are looking for an introduction to more advanced virtualization features at a highly affordable price.

- PowerVM Standard Edition

  Available on all IBM POWER5™ and IBM POWER6™ processor-based servers, it includes features designed to allow businesses to increase system utilization and the performance they require.

- PowerVM Enterprise Edition

  Offered exclusively on POWER6 processor-based servers, it includes all the PowerVM features.

Table 2-1 lists the PowerVM features available on each PowerVM edition.

Table 2-1 PowerVM features table

Features	**Licensed by**	**Processor support**
LPAR with dynamic reconfiguration	All editions (limited on Express Edition)	POWER4™, POWER5, POWER6
Capacity Upgrade on Demand	All editions (specific models)	POWER4, POWER5, POWER6 with IVM, HMC
Micro-Partitioning™	All editions	POWER5, POWER6
Shared Dedicated Capacity	All editions	POWER6
Multiple shared processor pools	Standard Edition, Enterprise Edition	POWER6 with HMC

Features	Licensed by	Processor support
Virtual I/O Server	All editions	POWER5, POWER6
Integrated Virtualization Manager (IVM)	All editions	POWER5, POWER6
Virtual SCSI	All editions	POWER5, POWER6
Virtual Ethernet	All editions	POWER5, POWER6
Live Partition Mobility	Enterprise Edition	POWER6
Simultaneous multithreading	All editions	POWER5, POWER6

2.1.1 Logical partitioning and processor capacity

PowerVM provides several technologies to increase server utilization and resource sharing.

Logical partitions

Logical partition (LPAR) technology has been available in IBM POWER®-based systems since 2001 with the introduction of the IBM pSeries® 690, a POWER4-based server. A set of system resources, such as whole POWER microprocessors, memory, and I/O resources, are logically grouped into a partition called an LPAR. Because resource allocation is a logical function, the amount of resources can vary according to need and availability within the physical server.

Shared processor LPAR, introduced with POWER5, allows enables you to partition a processor into virtual partitions and provides the flexibility to change the allocation of system resources dynamically for those environments. The Micro-Partitioning feature provides the capability to create multiple virtual partitions within a processor to a granularity of 1/100th of a processor, with a 1/10th of a processor partition minimum.

Any of the virtual servers can run on any of the physical processors. Thus, the processor resources are fully shared, making it possible to run the physical server at very high utilization levels.

Micro-partitions can be designated as either *capped* or *uncapped* modes:

- Capped

  A capped micro-partition has a defined processor entitled capacity, which it is guaranteed to receive needed. It not allowed to use more processing power than its assigned capacity (referred to as the *cap*).

- Uncapped

  An uncapped micro-partition is a logical partition that can use more processor power than its assigned processing capacity. The amount of processing capacity that an uncapped logical partition can use is limited only by the number of virtual processors assigned to the logical partition and the amount of unused processing capacity that is available in the shared processor pool. If multiple uncapped LPARs require additional processor capacity at the same time, the server can distribute the unused processing capacity to all uncapped LPARs. This distribution process is determined by the uncapped weight of each of the LPARs. Uncapped weight is a number (in the range of zero through 255) that the administrator sets for each uncapped partition in the shared processor pool. By setting the uncapped weight, any available unused capacity is distributed to contending LPARs in proportion to the established value of the uncapped weight.

Capped and uncapped micro-partitions can coexist and both receive the processor resources from the same physical shared processor pool.

Shared Dedicated Capacity

For POWER6 processor-based servers, the Shared Dedicated Capacity (SDC) feature allows for the *donation* of spare processor cycles from dedicated processor partitions to the shared pool, thus increasing overall system performance. The dedicated partition maintains absolute priority for dedicated processor cycles, and sharing occurs only when the dedicated partition has not consumed all its resources. Shared Dedicated Capacity is supported on IBM Power System POWER6 processor-based servers.

Multiple shared-processor pools

Multiple shared-processor pools (MSPP) allows for automatic nondisruptive balancing of processing power between partitions assigned to the shared pools, which results in increased throughput and the potential to reduce processor-based software licensing costs.

Figure 2-1 on page 21 shows an diagram of a POWER6 processor-based server with two processor pools defined.

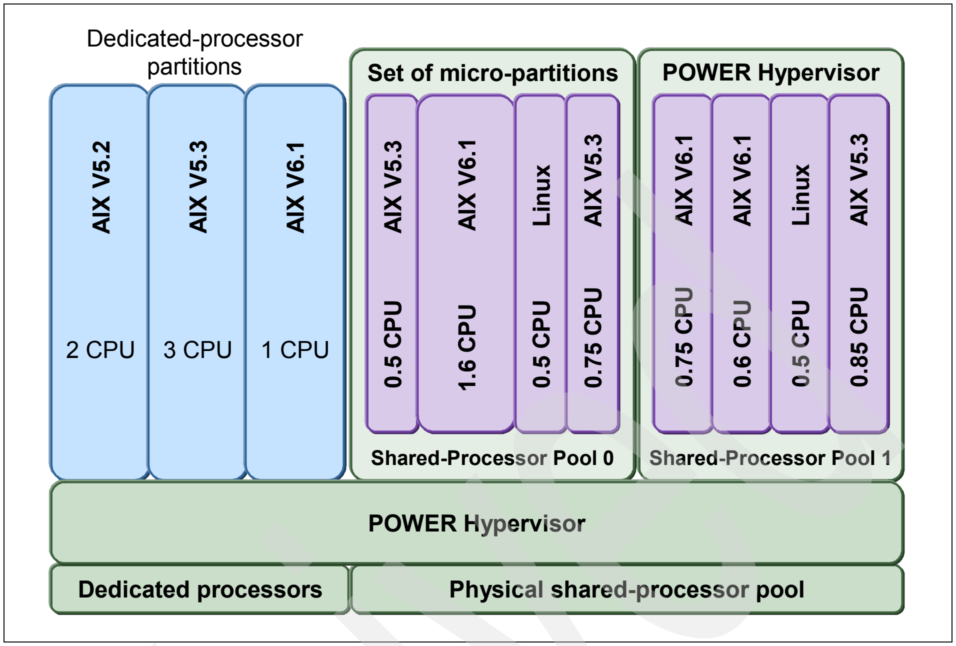

Figure 2-1 Architecture overview of POWER6 processor-based server with MSPPs

Partitions have a specific processing mode that determines the maximum processing capacity given to them from their shared-processor pool. The processing modes are:

- Capped mode

  The processing capacity given can never exceed the entitled capacity of the micro-partition.

- Uncapped mode

  The processing capacity can exceed the entitled capacity when resources are available in their shared-processor pool and the micro-partition is eligible to run. Extra capacity is distributed on a weighted basis. You must specify the uncapped weight of each micro-partition when it is created.

If competition exists for additional processing capacity among several uncapped micro-partitions, the POWER Hypervisor™ distributes unused processor capacity to the eligible micro-partitions in proportion to each micro-partition's uncapped weight. The higher the uncapped weight of a micro-partition, the more processing capacity the micro-partition will receive.

Simultaneous multi-threading

Simultaneous multi-threading (SMT) is an IBM microprocessor technology that allows two separate instruction streams (threads) to run concurrently on the same physical processor. SMT significantly improves overall processor and system throughput. SMT was first introduced on POWER5 servers and has been enhanced in POWER6 processor-based servers with gains in efficiency.

> **Note:** The benefit of SMT is greatest where there are numerous concurrently executing threads, as is typical in commercial environments, for example, for a Web server or database server. Some specific workloads, for example certain high performance computing workloads, generally perform better when SMT is disabled.

Capacity on Demand for IBM Power Systems

Several Capacity on Demand (CoD) possibilities are offered on selected Power System servers.

Table 2-2 lists the CoD options available for POWER5 and POWER6 processor-based systems.

Table 2-2 CoD features for POWER5 and POWER6

CoD Feature	Description
Capacity Upgrade on Demand	Enables permanent system upgrade by activating processors or memory
On/Off Capacity on Demand	Capacity-based billing, which allows for activation and deactivation of both processors and memory, as required
Trial Capacity on Demand	Partial or total activation of installed processors or memory for a fixed period of time

Table 2-3 lists the options available only to POWER5-based hardware.

Table 2-3 CoD options unique to POWER5 hardware

CoD Feature	Description
Reserve Capacity on Demand	Prepaid agreement that adds reserve processor capacity to the shared processor pool, which is used if the base shared pool capacity is exceeded

Table 2-4 lists the options available only to POWER6 processor-based hardware.

Table 2-4 CoD options unique to POWER6

CoD Feature	Description
Utility Capacity on Demand	Processor capacity, measured in processor-minutes, can be activated or deactivated. A prepaid or postpaid agreement is required.

2.1.2 Sharing resources

PowerVM virtualization has several features to increase the sharing and utilization of hardware. In this section, we introduce these features.

Virtual Ethernet

The virtual Ethernet function is provided by the POWER Hypervisor. The POWER Hypervisor implements the Ethernet transport mechanism an a virtual Ethernet switch that supports VLAN capability. Virtual Ethernet allows secure communication between LPARs without requiring a physical I/O adapter or cabling.

You can use the virtual Ethernet feature with or without the Virtual I/O Server.

The virtual Ethernet feature can only provide network services between the partitions hosted on a single physical machine. The two ways to connect a virtual Ethernet to an external network are:

- Routing

  You can use an LPAR to perform Layer-3 IP packet routing. The partition requires a dedicated Ethernet adapter, and acts as a router between both networks.

- Bridging

  You can use the Virtual I/O Server to act as a Layer 2 Ethernet frame bridge. This function is provided by the Shared Ethernet Adapter Virtual I/O Server function. Another way to implement bridging is to use a Linux partition and use the **`brctl`** Linux command to create a Layer 2 Ethernet bridge

Shared Ethernet Adapter

A Shared Ethernet Adapter (SEA) can be used to connect a physical Ethernet network to a virtual Ethernet network. It also provides the ability for several client partitions to share one physical adapter. With a SEA, you can connect internal and external VLANs by using a physical adapter. The SEA that is hosted in the

Virtual I/O Server is a Layer 2 network bridge to securely transport network traffic between virtual Ethernet networks and physical network adapters. It cannot be run in a general purpose AIX or Linux partition.

Virtual SCSI

The functionality for virtual SCSI is provided by the POWER Hypervisor. Virtual SCSI allows secure communications between partitions and a Virtual I/O Server that provides storage backing devices.

The combination of virtual SCSI and the Virtual I/O Server capabilities allows you to share storage adapter bandwidth and optionally to subdivide single large disks into smaller segments. The adapters and disks can then be shared across multiple partitions, increasing utilization.

The virtual disk I/O capability offered by the combination of virtual SCSI and the Virtual I/O Server provides the opportunity to share physical disk I/O adapters in a flexible and reliable manner. A single physical disk I/O adapter and associated disk subsystem can be used by many LPARs on the same server. This facilitates the consolidation of disk I/O resources and minimizes the number of disk I/O adapters required.

Virtual I/O Server

The Virtual I/O Server is a special-purpose partition, called the *hosting partition*, which provides virtual I/O resources to client partitions. The Virtual I/O Server takes control of the physical I/O resources that are shared with clients. A physical adapter that is assigned to a partition can be shared by one or more partitions, enabling administrators to minimize the number of physical adapters they require for individual clients.

The Virtual I/O Server is thus designed to reduce costs by eliminating the requirement for dedicated network adapters, disk adapters, and disk drives. Unlike other virtualization techniques, PowerVM does not require all devices to be virtualized. Therefore, you are allowed to have configuration scenarios in which you can access a mix of virtualized and dedicated physical I/O devices from an LPAR.

Devices used by the Virtual I/O Server can be:

- A mixture of dedicated devices assigned to partitions for maximum performance
- Devices that are used in the Virtual I/O Hosting Partition to be shared by multiple partitions, to provide higher efficiency of resources and adapters

Integrated Virtualization Manager

The Integrated Virtualization Manager (IVM) enables you to point, click, and consolidate workloads with its easy-to-use, browser-based interface. The IVM lowers the cost of entry into virtualization because it does not require the use of a hardware management console for single system partitioning.

With IVM, you can partition a single system, including the creation of LPARs and management of virtual storage and Ethernet. It is packaged as part of the Virtual I/O Server.

PowerVM Lx86

Run x86 Linux applications on POWER. This feature enables the dynamic execution of x86 Linux instructions by mapping them to instructions on a POWER processor-based system and caching the mapped instructions to optimize performance. PowerVM Lx86 software has features that enable users to easily install and run a wide range of x86 Linux applications on Power Systems platforms with a Linux on POWER operating system.

2.1.3 AIX 6 workload partitions

In AIX 6, workload partitions (or WPARs) add an additional operating system software based layer for virtualization of operating environments. Each workload partition can host applications and isolate them from applications executing within other WPARs. This capability can be leveraged on any server platform capable of running AIX 6, including POWER4, POWER5, POWER5+™, and POWER6 processor-based systems.

Figure 2-2 shows at what level the WPAR virtualization happens.

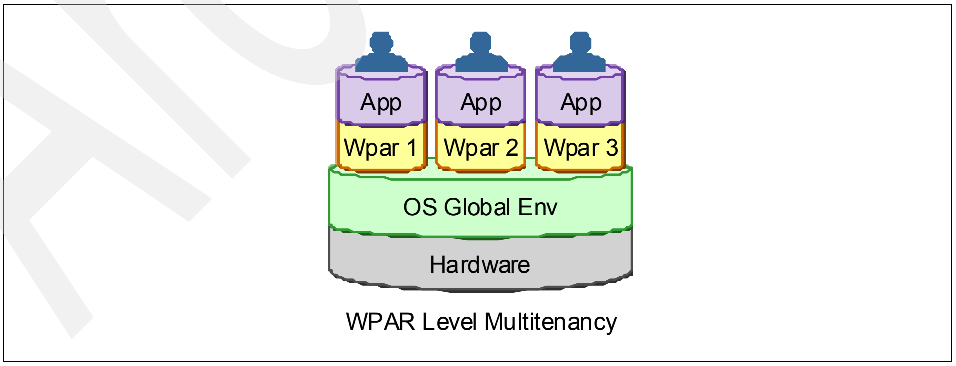

Figure 2-2 WPAR level multitenancy

Workload partitions can be created within an AIX 6 LPAR. Each workload partition provides an isolated environment for the application it hosts. From the application or service point of view, the WPAR provides an exact replica of a standard AIX operating system environment. Furthermore, the WPAR runtime environment can be dedicated to only hosting that application (the workload), and can be tuned to optimize performance, based on the specific workload characteristics of that application. Logically, WPARs can be considered as an operating system-level boundary around a specific set of AIX processes. Inside the WPAR, the applications have the following benefits:

- Private execution environments
- Granular control over processor and memory resources consumed
- Isolation from other processes outside the WPAR
- Dedicated network addresses and file systems
- Interprocess communication that is restricted to processes executing only in the same workload partition

Global environment

As mentioned earlier, workload partitions are created within standard AIX 6 instances, and the global environment is the part of an AIX 6 instance that does not belong to any workload partition. The global environment is therefore similar to the operating system environment of earlier versions of AIX. This global environment can be hosted within a dedicated LPAR or a micro-partition.

The global environment owns all physical or virtual resources of the LPAR. Resources are network adapters, disk adapters, disks, processors, and memory. The environment allocates processor and memory resources to the workload partitions, and provides them access to the network and storage devices.

The global environment has visibility into the workload partitions. A system administrator must be logged in to the global environment to create, activate, and manage workload partitions. Workload partitions cannot be created within other workload partitions. From the global environment, you can see (and control) the processes that are executing within the WPARs, and see the file systems used by the WPARs (non-root users cannot see WPAR data). For this reason, no user accounts, other than the system superuser, should have access to the global environment.

Performance monitoring and tuning activities can be enabled from the global environment and within the WPARs themselves.

Figure 2-3 on page 27 shows a diagram of three WPARs running specific applications.

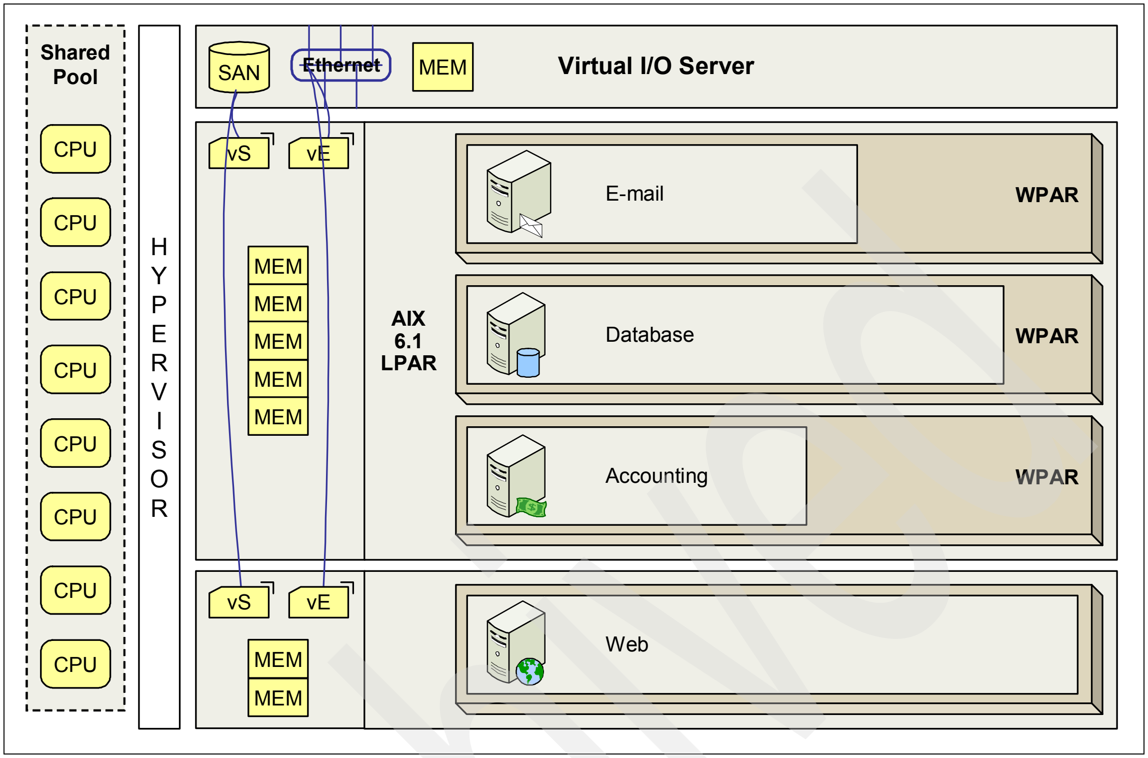

Figure 2-3 Three WPARs configured under one LPAR

System WPAR

A system WPAR is similar to a typical AIX environment. Each system WPAR can have dedicated writable file systems, but it also shares the global environment `/usr` and `/opt` file systems in read-only mode. This mode is the primary mode of operation for system WPARs. When a system WPAR is started, an init process is created for it, which in turn spawns other processes and daemons. For example, a system WPAR contains an inetd daemon to allow complete networking capacity, enabling the possibility to remotely log in to a system WPAR. It also runs a cron daemon, so you can schedule jobs.

Application WPAR

If an application or group of applications can be started by using a single command in the AIX command-line interface (such as a script), then the application can be hosted by an application WPAR. When the command exits, the workload partition is also automatically removed. Application WPARs are a quick way to leverage the isolation, resource control, and mobility features of workload partitions for hosting virtually any application or process.

An application WPAR shares the file system of the global environment and may also mount remote file systems.

An application WPAR can run daemons, but it does not run any of the system service daemons such as inetd, cron, or srcmstr. Remotely logging in to an application partition or remotely executing an action in an application WPAR are not possible. No kernel extension or device driver can be loaded into a WPAR. For example, GPFS file system can be mounted only in the global environment.

2.1.4 Partition mobility

Partition mobility allows dynamic consolidation of resources. For example, during non-peak usage hours, you can consolidate workloads on a fewer number of machines.

This powerful feature can be used to improve resource utilization, reduce power consumption, and increase serviceability of the POWER platform.

The two types of POWER partition mobility that we explore in this section are:

- WPAR Live Application Mobility is a feature of AIX 6 and WPAR Manager. It is available on POWER4, POWER5, and POWER6 processor-based systems.
- Live Partition Mobility relies on the POWER6 processor-based hardware and Hypervisor technology (Advance Power Virtualization). It is available on POWER6 processor-based systems only, but is supported on AIX 5.3, AIX 6, and Linux LPARs.

WPAR Live Application Mobility

Both types of workload partitions, the system WPAR and the application WPAR, are capable of being configured to support mobility, or *relocation*.

The capability to move one WPAR from one LPAR to another, possibly from one physical system to another, can be executed on active partitions. In this case, the application undergoes active relocation (it is *hot-migrated*) without stopping the application. The only visible effect for a user of the application is a slightly longer response time while the application is migrating.

WPAR Live Application Mobility is not a replacement for a high availability solution. The premise allows for planned migrations of workloads from one system to another so that the application is uninterrupted, for example, during hardware maintenance or a firmware installation on the server. The workload does not have to be aware of the migration for the most part. However, proper planning and testing are always recommended before moving anything into a production environment.

Note: Workload partition mobility is a software solution that depends on AIX 6 for execution. When used for the migration of a WPAR from one LPAR to another or between physical systems, then *hardware and software compatibility are required*.

Live Partition Mobility

As part of the PowerVM Enterprise Edition offering, Live Partition Mobility allows clients to move a running partition from one physical Power System POWER6 processor-based server to another POWER6 server without application downtime, helping clients to avoid application interruption for planned system maintenance, provisioning, and workload management. Live Partition Mobility is supported on IBM Power System POWER6 servers. Live Partition mobility also should not be considered a replacement for a high availability solution.

The migration operation, which takes a few seconds, maintains complete system transactional integrity. The migration transfers the entire system environment, including processor state, memory, attached virtual devices, and connected users.

Live Partition Mobility allows you to move partitions around, so you can perform otherwise disruptive maintenance operations without affecting partition availability. This approach can allow maintenance to be performed when it is convenient to you instead of when it is convenient to the users. Live Partition Mobility helps you meet increasingly stringent service level agreements (SLAs), because it allows you to move running partitions and applications proactively from one server to another.

The ability to move running partitions from one server to another offers the ability to balance workloads and resources. If a key application's resource requirements increase unexpectedly to a point where there is contention for server resources, you might move it to a larger server or move other, less critical, partitions to different servers and use the freed up resources to absorb the peak.

Live Partition Mobility can also be used as a mechanism for server consolidation, because it provides an easy path to move applications from individual, stand-alone servers to consolidation servers. If you have partitions with workloads that have widely fluctuating resource requirements over time (for example, with a peak workload at the end of the month or the end of the quarter) you can use Live Partition Mobility to consolidate partitions to a single server during the off-peak period, allowing you to power off unused servers. Then, move the partitions to their own, adequately configured servers, just prior to the peak. This approach also offers energy savings by reducing the power to run machines and the power to keep them cool during off-peak periods.

Security considerations for partition mobility

When moving a partition, the memory content of a running OS is transmitted over a network and onto another physical machine. This data contains both the user space and the kernel memory. Therefore, the security of the network must be considered. The sensitivity of any data or passwords that might be in memory during mobility must be considered because this memory is exposed in transport over the network. Additionally, if the network is not trusted and no encryption or Virtual Private Network (VPN) is in place, a theoretical possibility is that a man-in-the-middle attack can inject a security vulnerability into the executable sections of memory while in transit.

2.1.5 Reliability, availability, and serviceability

Because individual Power System servers are capable of hosting many system images, the importance of isolating and handling service interruptions becomes greater, especially in a multitenant environment. These service interruptions can be planned or unplanned. Interruptions for systems maintenance should be carefully considered, to ensure you consider the resources that will be affected by the downtime.

Technologies such as Live Partition Mobility and WPAR Live Application Mobility can be used to move workloads between machines, allowing for scheduled maintenance, minimizing any service interruptions.

IBM takes a holistic approach to systems reliability, availability and serviceability, from the microprocessor, which has dedicated circuitry and components designed into the chip, to Live Partition Mobility and the ability to move running partitions from one physical server to another. The extensive component, system, and software capabilities that focus on RAS, coupled with good systems management practice, can deliver near-continuous availability.

Note: For more detail about AIX RAS features, see *IBM AIX Continuous Availability Features*, REDP-4367.

2.2 PowerVM technology on resource sharing spectrum

Multitenancy is a way of realizing savings by sharing and consolidating resources, being more reactive to business changes, and providing services at lower costs. We explore where in the resource sharing spectrum some of the PowerVM technologies can be placed.

Table 2-5 on page 31 lists the abbreviations used in the figures in this section.

Table 2-5 PowerVM technologies abbreviations.

Abbreviation	PowerVM technology
WPAR	AIX 6 Workload Partition
VSCSI	Virtual SCSI
LPAR	Logical Partition
SLPAR	Share Processor LPAR
MSPP	Multiple Shared Processor Pool
SDC	Shared Dedicated Capacity

Figure 2-4 shows how all the PowerVM technologies help to increase resource sharing. Because IVM and partition mobility are not resource sharing technologies, they are excluded from the figure.

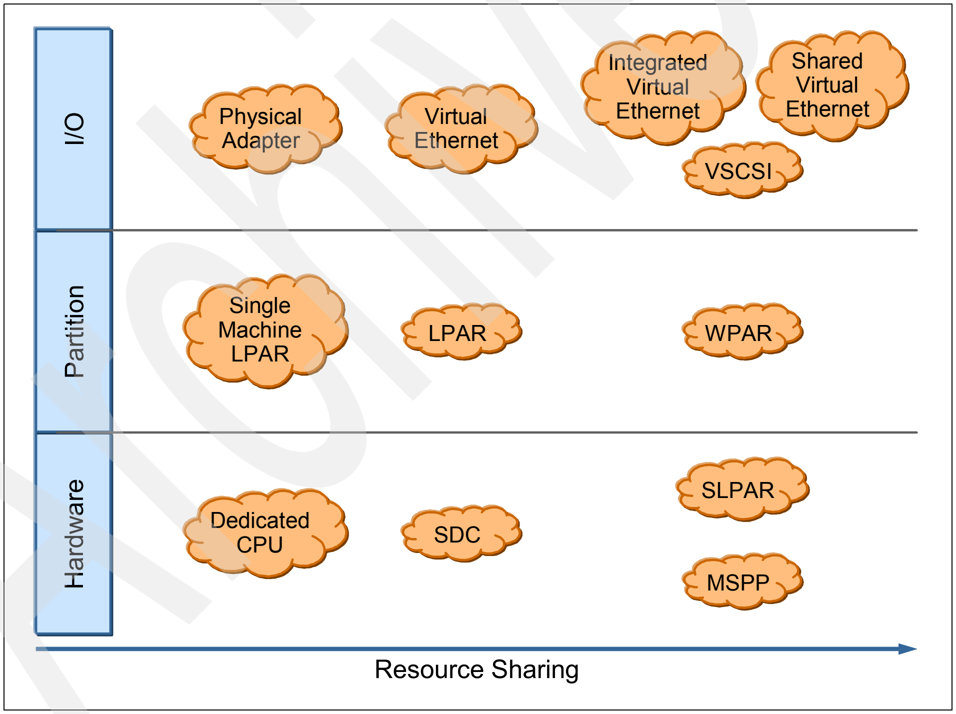

Figure 2-4 PowerVM technology on the spectrum of resource sharing

2.2.1 Multitenant case scenarios

For ease of discussion, we define four multitenant case scenarios (A-D) and discuss the best way of implementing multitenancy by using specific PowerVM technologies. Notice that we will try to pick some scenarios that will represent some tenant situations but not all of them. You can easily combine the concepts from several scenarios to create a more complex scenario, that would cover different profiles of tenants.

> **Note:** Because the audience of this book can be multitenant utility service providers, we use the term tenant instead of customer. And for a single customer, you could provide multiple tenant services. The *tenant* is the application and its required resources.

- Multitenancy case A

  In case A, we are trying to service tenants that demand a high level of resource isolation and the security benefits that come from that. Hardware isolation is very important for them and they are willing to pay a higher price to get it. They want full control of their OS instances and they require very high serviceability and uptime.

  Tenants usually run certain I/O intensive applications, with large databases. For example, banking or retail businesses both use large databases.

- Multitenancy case B

  In case B, our tenants do not mind sharing resources with other tenants, if a clearly defined boundary exists between tenants and cost savings are associated with it.

  Tenants will be running mission critical applications, which do not have heavy I/O, but the applications have to be available all the time. The tenants want full control over their resources.

  An example is tenants that are running customized Web applications, which are in the front-end for each tenant. Each tenant prefers to customize its OS, try different applications, and have the power to manage the OS.

- Multitenancy case C

  In case C, we have medium to small sized tenants that are very cost sensitive. They are usually divided in departments (development, test and support) that do not require strong security boundaries between them.

  The tenants run several business-critical applications that are distributed and the tenants are receptive to reducing their management costs.

  An example is a store chain that wants to run internal applications and Internet applications.

- Multitenancy case D

  In case D, we have multiple tenants that are not trustworthy. The service provider must ensure that tenants cannot cross tenant boundaries or try to access resources that they are not entitled to access. The service provider has to make sure that when a security breach occurs that it is contained to the minimal possible resources. The tenants are very cost-sensitive and want to lower their management cost. Their applications are all Internet applications and are subject to frequent hacker attacks. The tenants are open to use prepackaged applications with minimal customization.

  Basically, security and cost are the two driving factors for these tenants.

  An example is a Web hosting company, which hosts sites on the Internet with high visibility.

Based on the these multitenant scenarios, we will suggest multitenant architectures that use Power virtualization technologies.

> **Note:** We are not taking into consideration redundancy on the proposed solution. In a real world situation, redundancy is critical to any architecture.

Multitenancy case A

For this case, the tenants do not want a lot of resource sharing. So our architecture falls on the left of the resource sharing spectrum (see Figure 2-4 on page 31). We do not use WPAR, Virtual I/O Server, Virtual SCSI, Virtual and Shared Ethernet.

Figure 2-5 on page 34 presents a high-level diagram of the proposed architecture.

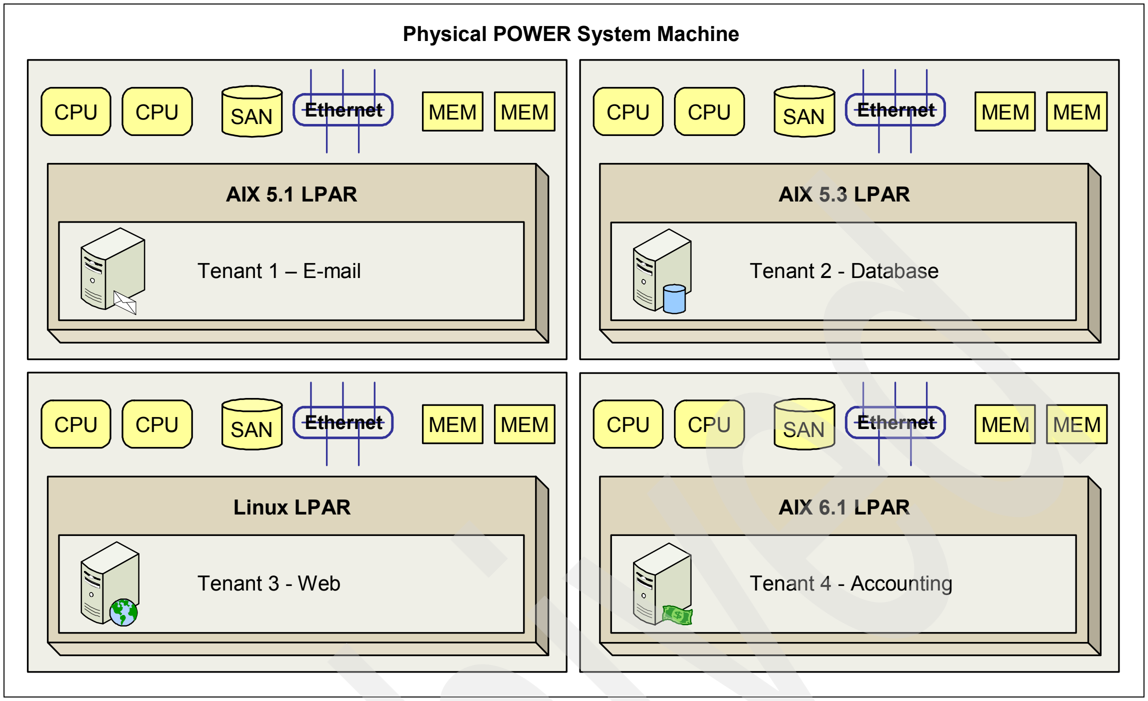

Figure 2-5 Multitenancy architecture for case A - using LPARs

The solution proposed hosts four tenants on the same machine, with individual LPARs. Each LPAR has dedicated CPU, memory, storage, and Ethernet adapter.

The following technologies can be used without affecting the customer perception of resource sharing:

- Shared Dedicated Capacity

  This feature can be used to fully utilize CPU cycles that are allocated but not being used by other LPARs. If one LPAR is requiring lots of CPU, space unused capacity can be allocated to the requesting LPAR.

  > **Note:** Certain considerations apply to the way processors can be shared. See *PowerVM Virtualization on IBM system p: Introduction and Configuration Fourth Edition,* SG24-7940.

- Shared-processor LPAR (SLPAR)

  Instead of using dedicated CPUs, the service provider can create SLPARs, that are assigned a share of the processor. The share can be dynamically changed.

- Multiple shared-processor pools (MSPP)

  This feature can be used to concentrate CPU resources to a set of LPARs. One example is to make a CPU pool that is used by the production and development databases. The production is an uncapped CPU partition, so more CPU resources can move to it when they are required.

- Dynamic LPAR

  Dynamically increasing or decreasing the number of processors, memory, and adapter slots assigned to a LPAR is possible. One consideration when using dynamic LPARs is that the HMC must have TCP/IP connectivity to the LPAR.

  For details of this capability, see *PowerVM Virtualization on IBM system p: Introduction and Configuration Fourth Edition,* SG24-7940.

Each tenant is on a separate TCP/IP network and has full control over its LPAR. For that reason certain security aspects have to be managed by an acceptable usage policy, such as:

- Changing IP addresses or host name

  The service provider has to provide a policy regarding IP address changes, because the change can affect availability monitoring and billing.

- Powering on and off

  This policy can affect the service provider availability monitoring or SLA terms.

- Usage data collection

  The service provider has to be granted access to the OS instance to collect usage information, if the billing is more granular than a single LPAR.

Multitenancy case B

For this case, we have tenants that can use a higher degree of resource sharing. Figure 2-6 on page 36 presents the architecture.

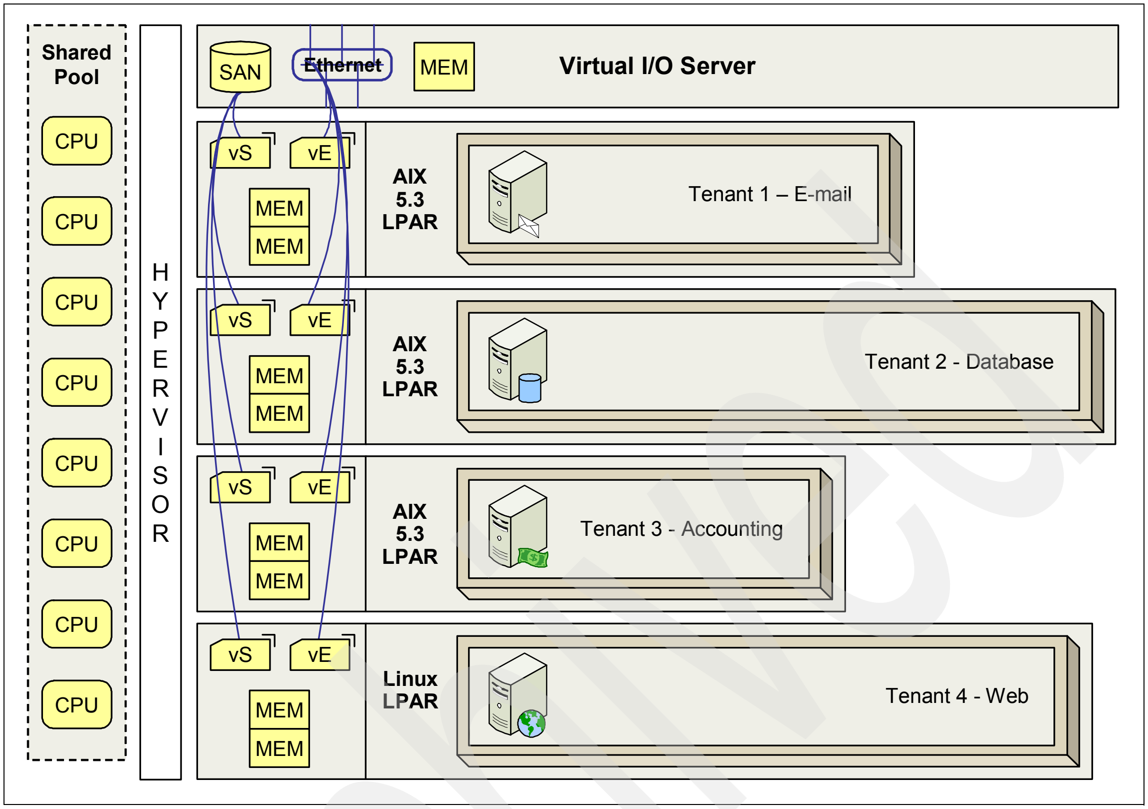

Figure 2-6 Multitenancy architecture for case B - tenants sharing resources

We use the following PowerVM virtualization features:

- Micro-Partitioning
- Virtual I/O Server
- Virtual Ethernet, Shared Ethernet
- Virtual SCSI
- Multiple shared-processor pools

In this architecture, we have LPARs that share a single processor pool. For Tenants 2 and 4 (database and Web tenants), their LPARs can be configured to receive extra processor cycles when necessary.

All network and disk I/O is virtualized through the Virtual I/O Server. The reliability of the configuration can be increased by adding a second Virtual I/O Server to create a multipath for the storage and network. For sensitive data that might reside on a storage area network (SAN) or any other type of storage, consider using encryption solutions such as AIX Encrypted File Systems (EFS).

EFS can ensure that sensitive data is protected with strong encryption across the network, through the Virtual I/O Server and is only decrypted after it accesses system memory again.

Live Partition Mobility can be used to distribute load or move workloads when performing disruptive maintenance on a machine is necessary.

Each tenant has full access to its LPAR and can manage it accordingly to the tenant's policies.

In this configuration, certain considerations have to be managed by an acceptable usage policy, such as:

- Changing IP addresses or host names

  The service provider has to provide an acceptable policy regarding IP address changes, because this can impact availability monitoring and billing.

- Powering on and off

  This can affect the ability of the service provider to monitor the machine availability.

- Usage data collection

  If the billing is more granular than a single LPAR, the service provider has to be granted access to the OS instance to collect usage information.

Multitenancy case C

For this case, we use the same PowerVM virtualization features, but we increase the level of resource sharing by using WPARs.

Figure 2-7 on page 38 depicts the new configuration. WPAR features are used in different ways for Tenant 1 and Tenant 2:

- Tenant 1

  Tenant 1 has control over the global environment WPAR. This approach means that the tenant is able to control the OS level, what applications are installed, and changes to the OS configuration. This configuration is ideal for development, test, and production environments of an application. Each environment is identical, but is logically separated from each other, which allows for the testing, development, and production on the same environment.

  Another possible usage is for a company with tightly coupled departments. Each department is assigned an individual WPAR that they can use for their applications. The IT Department controls the global environment and ensures that everyone is at the same OS and application level. The individual departments have full control over their WPARs.

- Tenant 2

  In the Tenant 2 environment, the global environment is controlled by the service provider. This approach allows the service provider to control what OS level, applications, OS configuration, and file systems that are available to the additional WPARs. Tenant 2 has access to all local resources of the local WPAR. This configuration is ideal for a utility application provider, by permitting greater control over the execution environment.

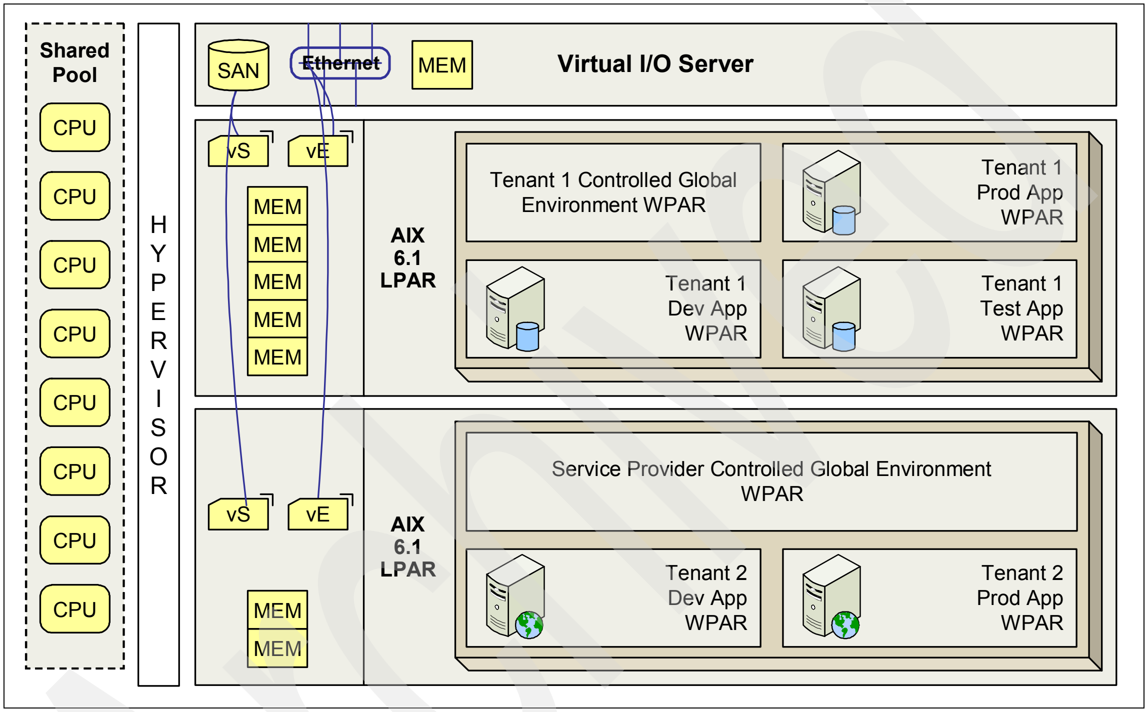

Figure 2-7 Multitenancy architecture for case C - using WPAR

Each tenant uses different features of WPAR. Tenant 1 could also use the application WPAR to isolate a single application.

Tenant 2 is giving away some control of the environment. The service provider could use this configuration to offer a standard, out-of-the-box solution for Web hosting, for example, where the tenants have limited control over they Web instances.

The Tenant 2 configuration provides a greater degree of security isolation. Compared to the solutions for cases A and B, the Tenant 2 configuration provides a greater degree of control for the service provider. There is a reduced set of configuration that the tenant can perform inside the WPAR environment. The tenant will not be able to change IP addresses or Ethernet MAC address, or be able to upgrade the OS.

The Tenant 1 configuration provides more flexibility to the tenant. The service provider still has the same security level as on the multitenant case B, but the tenant has a greater deal of control over subtenant's WPAR. This might make the solution more valuable to the tenant, compared to the single LPAR solutions.

Multitenancy case D

For this case, the main concern is security, followed by cost. The proposed solution, illustrated on Figure 2-8 emphasizes security and a high level of resource sharing.

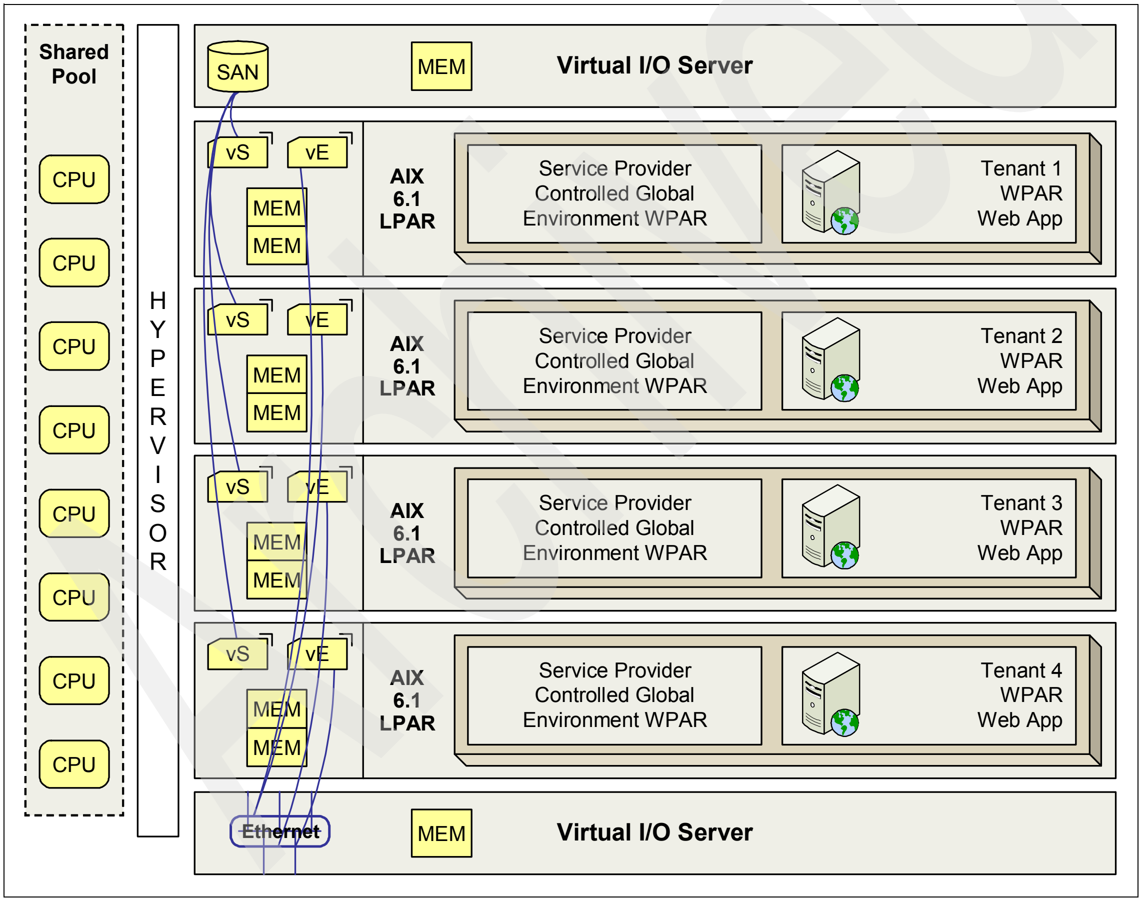

Figure 2-8 Multitenancy architecture for case D - using WPAR to increase security isolation

Although this solution uses the same features as multitenancy case B, in this case we use the AIX 6 WPAR feature to add an extra layer of system isolation to each tenant. Any breach of security is limited to the tenant WPAR. If a buffer overflow is found on the Web server, only the WPAR will be compromised. This creates an additional security perimeter around the compromised application.

The solution proposes using two Virtual I/O Servers, one for the Shared Virtual Ethernet and the second for the Virtual SCSI. This approach can isolate any possible denial-of-service attack in the network or on the disk I/O.

Note: The guidelines in the following Web site can help you configure the Virtual I/O Server to increase resilience to denial-of-service attacks:

http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/perf.html

Also, check the section 2.3.1, "Information technology security" on page 41 for a more detailed discussion about security aspects of this solution.

For the service provider, this solution presents the advantage of more control over the tenant operations. The majority of configuration, upgrading, and management can be made only on the global environment.

Note: For more details about WPAR configuration and management, see *Introduction to workload Partition Management in IBM AIX Version 6.1,* SG24-7431.

The service provider can also reduce costs by automating the provisioning, deployment and management of the WPAR instances. Usage and monitoring can be done from the global environment WPAR.

A unique feature of this solution is the possibility of having an Encrypted File System on the tenant WPARs that cannot be accessible to the service provider. That can be used to safeguard highly sensitive information.

Note: For more information about the Encrypted File System on a WPAR, also see *Workload Partition Management in IBM AIX Version 6.1,* SG24-7656.

This solution provides the highest degree of isolation and at the same time provides a great degree of data sharing.

Notice that the conventional wisdom says that as we increase resource sharing the security isolation will decrease. So, looking back at our proposed solutions the security isolation should decrease as we increase resource sharing, as shown in Figure 2-9.

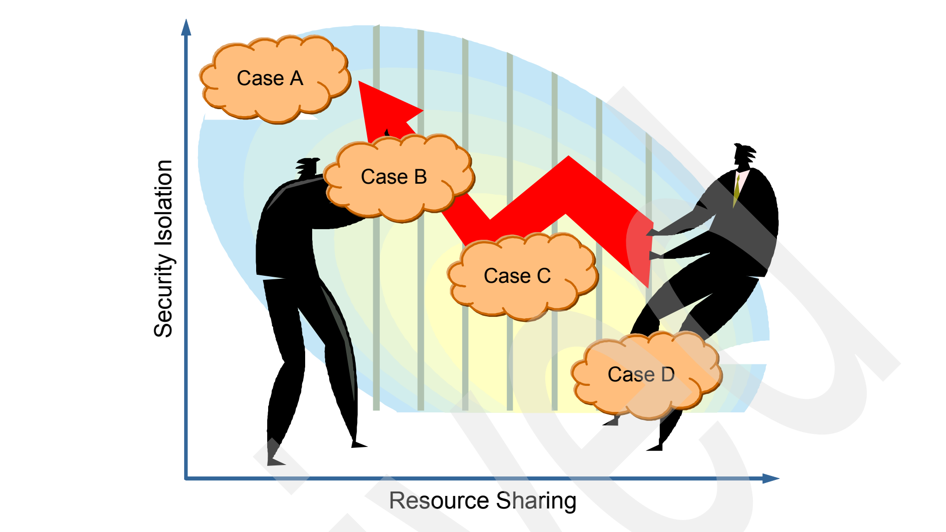

Figure 2-9 Conventional wisdom - a trade-off between security and sharing resources

However, security is not directly related to the resource sharing level, but also related to specific security features, functions, policies of the proposed architecture.

In the next section we provide a more detailed security analysis of each proposed scenario solution.

2.3 Additional considerations

In this section, we discuss several considerations that arise on multitenant environments. When moving from an isolated resources scenario to a shared resource scenario, new interactions can happen between the shared components, creating new concerns that should be addressed.

2.3.1 Information technology security

Information Technology Security is one of the most important aspects when considering a multitenant architecture. The architecture has to be designed from the ground up to implement a clear and well-defined boundary between tenants.

We use the following key concepts to analyze the security of multitenancy:

- Confidentiality

  Confidentiality means preventing unauthorized information disclosure. Only entities that are entitled to access the information should be able to do so. For example, when performing a credit card operation on the Internet, there are several safeguards to ensure the credit information is confidential:

  - The physical credit card should be kept in a secure place.
  - The credit card information should be encrypted when sent over to the merchant.
  - The credit card information should be protected when it resides in a merchant's database.
  - The credit card information has to be destroyed when it is not needed anymore.

  When the information is accessible to an unauthorized party, we say there was a *breach of confidentiality*. A breach of confidentiality can take many forms, such as someone looking over your shoulder, a stolen wallet, somebody *sniffing* the network communication, or a credit card receipt that is publicly available.

- Integrity

  Information integrity means that only authorized entities can modify the data. A breach of integrity happens when someone, who is not authorized, changes the data. For example, an unauthorized individual might be someone who sends a computer virus to erase information, or a rogue employee who modifies salaries. It could also be a computer administrator who changes files that should not be changed.

  Technical controls that are used to maintain information integrity are usually:

  - Access controls: These are user IDs, password, and tokens combined with access controls lists.
  - Security zones: These changes can only be made after changing the security privileges.
  - Privileges: Only specific entities with certain privileges can make changes.

- Availability

  Any information has to be available to be useful. This means availability of data, and the methods to ensure the data's security.

  Denials of availability can be, for example, hardware theft and denial-of-service (DoS) attacks.

General considerations

When we talk about sharing resources with multiple tenants, the first concern expressed is regarding security isolation between tenants. Indeed, multitenancy does create the following new challenges:

- Denial-of-service (DoS) attacks on shared resources

  Multitenant service providers can realize savings by sharing resources between tenants. One undesired side effect is that shared resources can be subjected to intentional or accidental denial of service, when one tenant starts using too much of the shared resource.

  There are too many vectors for DoS attacks, so the first defense is to have an effective resource usage monitoring. When properly configured, as described in the following list, Power Virtualization technologies can help mitigate the effects of DoS attacks:

  - CPU sharing

    LPAR can be configured in several ways, where only the spare CPU cycles are shared, so each LPAR has an assured minimum amount of spare processor capacity.

    For more details, see:

    - *PowerVM Virtualization on IBM system p: Introduction and Configuration Fourth Edition,* SG24-7940
    - *PowerVM Virtualization on IBM System p: Managing and Monitoring,* SG24-7590.

    For a background technical discussion about security isolation features of LPARs, see *Logical Partition Security in the IBM eServer™ pSeries 690*, available at the following location:

    http://www.ibm.com/systems/p/hardware/whitepapers/lpar_security.html

  - Network bandwidth sharing

    DoS attacks on network resources originating on one tenant can be disruptive for all tenants that share the same network connection.

    A service provider can use the AIX TCP/IP Quality of Service (QoS) feature combined with WPAR and a network switch that supports QoS to shape the traffic dynamically. When one tenant starts using too much traffic the service provider can tag all the traffic originating from that tenant to a lower QoS to guarantee that resources are available to all customers.

    This solution assumes that the service provider is using WPARs and retains administrator access to the global environment.

  - Disk storage sharing

    An I/O DoS attack is possible when you use the Virtual I/O Server. The following link has best practices for configuring the Virtual I/O Server to avoid attack situations:

    http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/perf.html

- Monoculture

  Although monoculture can ease the burden of management, by creating homogeneity of hardware, configuration, software installation, it can also amplify any security vulnerability. The service provider must always deploy safeguards to mitigate that risk.

- OS management privileges

  One common challenge in a multitenant architecture is how to manage the administrative boundary between the tenant and the service provider. Allowing the tenant to configure the OS can result in increased security risks and consequential impacts. The service provider has to retain the administrative privilege for tasks that can disrupt the infrastructure, like changing network settings, storage allocation, usage data, and logging. The tenant should retain control of the application. AIX 6 provides several features to accommodate separation of privileges, the main one being WPARs.

- Management and monitoring network

  Multitenant architecture usually shares a common management and monitoring network across multiple tenants. That architecture allows the service provider to use the same management and monitoring infrastructure for all tenants. One concern of this feature is that this network can be used as a vector of attack, because the network now provides a connection between all tenants. To mitigate this risk, the service provider should provide a firewall for all tenant connections to the management and monitoring network and allow only outbound initiated traffic to happen.

Security analysis of tenant scenarios

Because we can have several degrees of resource sharing on multitenancy, in the following sections, we analyze the security of each of the tenant scenarios and solutions that we defined in 2.2.1, “Multitenant case scenarios” on page 32.

Multitenancy case A scenario

In case A, we have a limited level of resource sharing and an excellent level of security isolation. In multitenancy case A, the PowerVM Hypervisor is the key technology enforcing access control, information integrity, and resource availability.

The PowerVM has a long history of reliability and security, as shown in Figure 2-10.

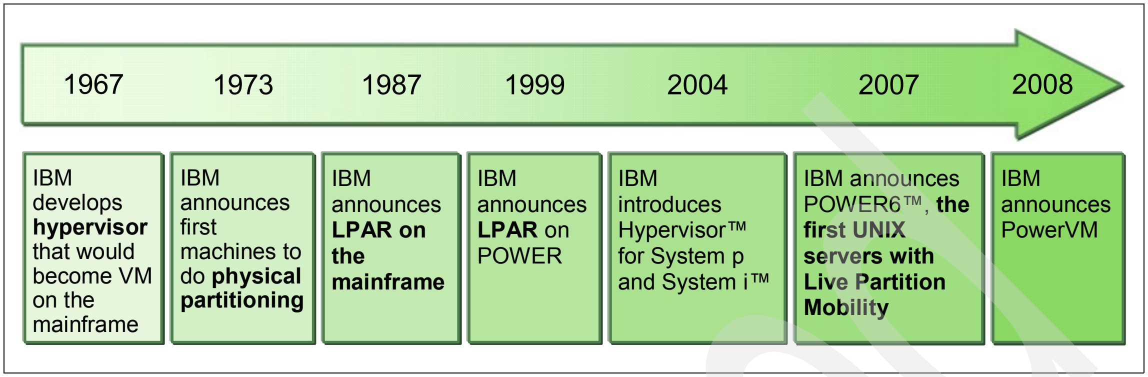
1967
1973
1987
1999
2004
2007
2008
IBM develops **hypervisor** that would become VM on the mainframe
IBM announces first machines to do **physical partitioning**
IBM announces **LPAR on the mainframe**
IBM announces **LPAR** on POWER
IBM introduces Hypervisor™ for System p and System i™
IBM announces POWER6™, **the first UNIX servers with Live Partition Mobility**
IBM announces PowerVM

Figure 2-10 IBM PowerVM Hypervisor history

In addition to the traditional *user mode* and *supervisor* mode processor states, the POWER5 and POWER6 processors have a new privileged processor state called *Hypervisor mode* that has been added to support logical partitioning. Hypervisor mode is accessed through Hypervisor call functions, which are generated by the operating system kernel that is running in the partition. The Hypervisor mode allows for a secure mode of operation that is required for various system functions where LPAR integrity and security are required. The Hypervisor validates that the partition has ownership of the resources it is attempting to access, such as processor, memory, and I/O, then completes the function. This mechanism allows for complete isolation of partition resources.

The IBM Power Hypervisor is a proven technology and has received several security certifications. For more information about the security certifications, see:

http://www.ibm.com/systems/p/os/aix/certifications/index.html

We can better understand the scenario security by looking at the three security concepts:

- Confidentiality

  Confidentiality is guaranteed by the OS and the hardware layer. The PowerVM Hypervisor controls the CPU and memory sharing, and the device assignment to specific partitions. Hypervisor assures that no information leakage happens between partitions and makes sure that access control to the resources are honored.

- Integrity

  The OS implements all the integrity controls, such as login access, file access permissions, and network configuration the same way as in single machine configurations.

  The IBM Hypervisor enforces the assignment of shared resources, so one LPAR cannot access resources that it does not own.

- Availability

  IBM PowerVM Hypervisor also implements controls to guarantee that an LPAR has all the resources that were allocated to the partition, preventing DoS attacks on shared resources. Several ways of configuring exist for the CPU sharing on the Hypervisor, so in case A we would probably use dedicated CPUs optionally with shared dedicated capacity. However, the service provider could also use the multiple shared processor pool functionality that provides a more flexible mechanism to allocated CPU across multiple LPARs, while providing a way to have guaranteed resources available.

Multitenancy case B scenario

Multitenancy case B builds on case A and adds Virtual I/O Server functionality to virtualize storage and network interfaces. As we already discussed the security implications of case A, this case B discussion is centered on the added Virtual I/O Server functionality.

We can better understand the scenario security by looking at the three security concepts:

- Confidentiality

  On the network side the Virtual I/O Server implements a virtual network switch supporting IEEE 802.1Q VLAN tagging. If you select the IEEE 802.1Q compatible adapter check box when you create the virtual adapter, you ensure that every data frame is tagged and only presented for the virtual Ethernet adapter in the same VLAN.

  Combined with a VLAN-aware physical network switch, you can isolate or group the LPAR network interfaces on a different VLAN, while at the same time, assuring that all traffic is protected from unauthorized access.

  On the storage side, the Virtual I/O Server acts as a SAN switch, so assigning exclusively to the partitions that will be using it is important.

  Figure 2-7 on page 38 shows the create virtual SCSI dialog and the option that should be selected.

  When the storage is assigned to a single partition, the Virtual I/O Server ensures that only that partition can access the storage.

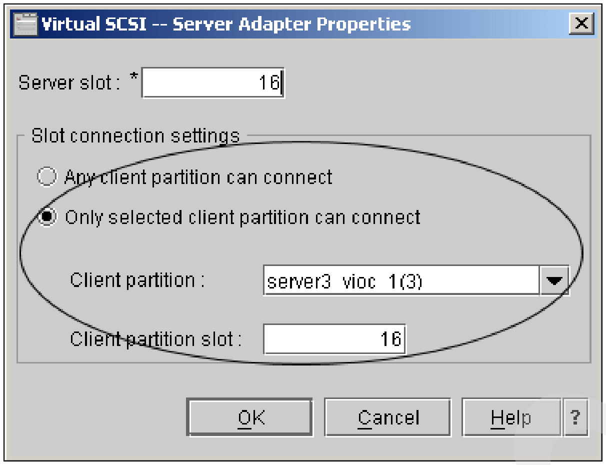
Virtual SCSI -- Server Adapter Properties

Server slot : * 16

Slot connection settings

○ Any client partition can connect

● Only selected client partition can connect

Client partition : server3 vioc 1(3)

Client partition slot : 16

Figure 2-11 Virtual I/O Servers SCSI client selection

- Integrity

  In the same way that the Virtual I/O Server protects information confidentiality, it protects integrity. On the network side, the Virtual I/O Server enforces network access by using VLAN tagging. On the storage side, only LPARs that were granted access to storage are able to access and modify it.

- Availability

  A DoS attack is possible when you use the Virtual I/O Server. The following link has best practices for configuring the Virtual I/O Server to avoid attack situations:

  http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/perf.html

Multitenancy case C scenario

Multitenancy case C provides a more complex configuration, making use of the Virtual I/O Server functionality and AIX 6 WPARs. As we already discussed, the Virtual I/O Server security considerations we concentrate on are WPAR features.

> **Note:** This section provides a introduction to WPAR security. See *Workload Partition Management in IBM AIX Version 6.1,* SG24-7656 for a more detailed look at security isolation provided by AIX 6 WPARs.

To understand WPAR security isolation, consider that each WPAR represents an individual environment, which relies on and uses services provided by the global environment.

The global environment represents a framework that controls all user-level and system-level objects that are usual for any regular AIX operating system image such as devices, file systems or processes. At the core of the whole WPAR technology is the shared kernel, meaning that the global environment and all active WPAR instances share a single kernel.

WPARs and global environment have been designed so that administrative tasks and commands that can be run from the global environment have the ability to affect WPAR operations, whereas the potential of a partition to interfere with a different partition or the global environment is strictly limited. The result is a set of WPAR considerations and customizations that includes the following categories:

- Device access
  - Access to storage devices, physical disk devices, and objects associated with Logical Volume Manager (LVM), such as physical volumes, volume groups or logical volumes are not available within a WPAR. Access to data is performed through file systems that are mounted from the global environment into the WPAR.
  - Physical devices such as network devices are not available in a WPAR. WPARs have access only to a set of pseudo devices such as `/dev/ptyp0`. Devices that could provide a more global view of the system such as `/dev/mem` or `/dev/kmem` have been removed, but access to system-generic devices that are safe, such as `/dev/null` or `/dev/zero` are allowed. WPARs are not capable of creating new devices by themselves, for instance when you use the **`mknod`** command.
- File system access
  - File systems within a system WPAR can only be accessed from that WPAR or from the global environment. File system objects that belong to the global environment and are then exported may be accessed within the normal constraints of shared access for those objects from both environments.
  - Although a physical file system can be mounted in a WPAR, certain device interfaces associated to the file system do not appear within the WPAR environment. This means that certain utilities that access file system metadata and require access to certain `/dev` devices do not work when they are invoked from the WPAR. For example, extending or defragmenting a file system is not allowed from the WPAR.
  - Regular global users cannot access WPAR file systems unless explicitly given specific privileges.
  - The WPAR administrator cannot mount or unmount WPAR file systems that are unmounted or mounted from the global environment.

- Network access
  - Modification of network settings, such as the IP address, is not allowed from the WPAR. Also, processes from a WPAR can bind, send, or receive packets only from the IP addresses and ports configured in the WPAR. The source address of outgoing IP packets must match the address assigned to the interface by the global environment.
  - Application processes are not allowed to bypass the network stack.
- System settings
  - Access to system-wide settings and objects are restricted.
  - Access to certain system-wide configuration files such as those contained in `/etc` has been restricted.
- Command line interface
  - The ability to use certain administrative commands that could affect the global environment has been removed. For instance, modifying the date or time from a WPAR, or binding a process to a specific processor is not allowed.
  - System-wide parameters cannot be changed from a WPAR by using the command line interface.
- Security and user management
  - User, group and security repositories of all WPARs are distinct entities and are different from the global repository.
  - Applications running in a WPAR derive their credentials from the local WPAR repositories. The credentials scope is confined to the WPAR.
  - The scope of WPAR root privileges is contained within the WPAR boundaries and cannot interfere with the global environment.
- Process resources and intercommunication
  - Processes that run in the global environment and have appropriate privileges can view and signal processes within a WPAR.
  - Processes belonging to a WPAR cannot be reassigned to a different WPAR.
  - Processes within a WPAR can only create, view and access the resources owned by the WPAR. These resources include non-shared file systems, devices, network interfaces, or ports. They can only access files located in the file systems mounted in the same WPAR.

  - Processes within a WPAR can only see processes running in the same WPAR, can send signals only to processes in the same WPAR, and can share System V IPC mechanisms (shared memory areas, message queues, semaphores) or POSIX IPCs only with other processes executing in the same WPAR.
  - Resources that are used by applications running in a WPAR are tagged with the ID of that WPAR.

- Kernel manipulation

  The ability to load or unload system-level device drivers and kernel extensions from a WPAR has been removed.

As the previous list of categories illustrates, WPARs can be used to add another layer of security isolation. The security isolation can be used to increase the overall security level or to create another management boundary, if a tenant IT department could control the global environment and each WPAR could be used by a different tenant internal department.

In the proposed architecture for multitenancy case C, as shown in Figure 2-7 on page 38, you can see the administrative boundary being used on Tenant 1. Tenant 1 controls the global environment and can offer individual WPARs to each internal department. This approach allows for a centralized control over the environment that each WPAR is using, while at the same time delegating some of the control to each individual WPAR administrator.

For Tenant 2, the service provider retained access to the global environment. For this case, we assume that main reason for the tenant’s decision is to outsource the OS management, instead of having to increase the security isolation.

We can better understand the scenario security, by looking at the three security concepts:

- Confidentiality

  WPAR can increase the security isolation of data. For Tenant 1, we have another degree of security separation, where system administration privileges are separated in layers and assigned to different people. This situation tends to increase the overall security, by creating a separation of duties and avoiding administrators with too much privileges.

  WPARs administrators can use the AIX 6 Encrypted File System (EFS) to restrict access to data housed on WPAR. The administrator of the global environment would not have access to data on the WPAR.

> **Note:** See *Workload Partition Management in IBM AIX Version 6.1,* SG24-7656 for a detailed example of how to use the EFS inside WPARs.

- Integrity

  By providing different security realms, WPARs can maintain information integrity and protect information access. Each individual WPAR can have a separate user ID set, independent from the global environment.

- Availability

  Even though all WPARs share the same kernel, several parameters can be configured to prevent a WPAR from accidentally or intentionally exhausting system resources and affecting the performance of the other WPARs that are running under the same global environment. The WPAR resource control is based on the Workload Manager (WLM) technology, which has been incorporated on the AIX kernel since version 4.3.3.

  The following parameters can be configured to control resource usage:

  - CPU and memory

    The two methods to configure CPU and memory sharing are:

    - Shared, in which all resources are shared in an equal share between all the active WPARs running.
    - Percent-based, in which each WPAR is assigned a percentage of the resources. There are three values that can be assigned:

      - Minimum percentage: the minimum amount each WPAR will receive at all times.
      - Soft maximum: a maximum that can be exceed if there are no other contenders for the resource.
      - Hard maximum: a value representing the maximum amount of resources that can be used. That value cannot be exceeded.

  - Processes and threads

    Each WPAR can have a maximum number of processes and threads.

  - Virtual Memory

    Each WPAR can have a maximum amount of virtual memory configured.

  - Resource (CPU) set

    A resource set is used to define a subset of processors in the system. If a resource set is specified for a WPAR it may use only the processors specified within the resource set.

    > **Note:** Resources sets can also be used to achieve processor affinity on computing-intensive applications.

The WPARs feature of AIX 6 can greatly increase the security of architectures that make use of this feature.

Multitenancy case D scenario

In multitenancy case D, we added three new concepts:

- Using two Virtual I/O Servers to increase tolerance to DoS attacks

  As we discussed in “General considerations” on page 43, the Virtual I/O Server can be the target of DoS attacks. By splitting the network and storage functions across two Virtual I/O Servers, you can fine-tune each Virtual I/O Server to better tolerate these situations.

- Using WPARs to increase security isolation of tenants

  WPARs provide an additional level of security isolation, where the service provider has control over operations that are typically under the control of tenants. That additional privilege allows the service provider to have a greater control over the operations that are being performed by the tenants, thus reducing security exposure.

- Using individual WPARs to increase tenant isolation

  Although WPARs can be used to increase resource sharing, in this case we used WPAR to create an additional security defense perimeter around the tenant. WPARs provided another set of technical controls that will have to be crossed before a tenant can compromise another tenant security.

These three security features are added on top of the other features already discussed in cases A, B, and C.

To understand the scenario security, consider the three security concepts:

- Confidentiality

  Information confidentiality is maintained at several levels:

  - WPAR: file system controls, user ID controls, file controls
  - AIX 6 OS: file system controls, user ID controls, file controls
  - Virtual I/O Server: Virtual SCSI controls, VLAN tagging
  - Hypervisor: CPU, Memory and devices controls

  Each technology layer implements extensive access controls to guarantee that only authorized entities can access the data.

- Integrity

  The same technologies (listed for confidentiality) also ensure that only authorized entities may change the data.

Another feature of the layered technology is the separation of duties. Different groups can manage different levels of the infrastructure, guaranteeing an excellent level of separation of duties.

Duty separation is important for avoiding security leaks from inside a company, usually done by employees with administrator access.

- Availability

  This case D architecture uses two Virtual I/O Servers, one for the network and one for the storage, to increase the resilience of the architecture against DoS attacks in the network or storage resources.

  > **Note:** Use the guidelines in the following link when you configure your Virtual I/O Server to increase resilience to DoS attacks:
  >
  > `http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/perf.html`

The WPAR AIX 6 technology can be used as a tool to greatly increase the security level of architectures that use the Power Systems.

2.3.2 Resource starvation

A classical problem on shared resources systems is called *resource starvation*. As we increase resource sharing on a multitenant architecture, we also increase the risk that we will starve a resource, can cause either a deadlock or a serious performance degradation.

Resource starvation is usually the way the DoS attacks work. By exhausting one key specific resource, the DoS attack can render the whole system unresponsive or seriously degrade performance for all tenants.

When sizing a typical architecture without multitenancy, we usually size the resources for the worst case, having space capacity to handle peak loads. On a multitenant architecture, the sizing should be a little higher than the average load. By pooling resource capacity from each individual tenant we are able to handle individual peak loads. As a result, we cannot have service peak loads for all clients at the same time.

The mechanism to avoid resource starvation is highly dependent on the architecture and type of service being used. However, several ways to avoid resource starvation include:

- Proper capacity planning

  Resources should be sized to be able to accommodate the projected load. See the following publication for more details: *IBM eServer pSeries Sizing and Capacity Planning: A Practical Guide*, SG24-7071

- Distributed architecture

  To avoid resource starvation or bottleneck, the infrastructure must be designed to easily expand capacity if needed. All elements of the architecture should favor distributed architectures, instead of centralized one. Examples are load balancers for Web traffic, distributed storage, and others.

- Quality of service (QoS_

  QoS is a vital way of reducing resource starvation. QoS provides a way to arbitrate resource allocation, where high priority requests can be fulfilled first instead of in the order they were received.

- Resource throttling

  One way of avoiding resource starvation is to throttle the request rate. As the system approaches the resource exhaustion threshold limit, the resource fulfillment rate can be reduced to avoid starvation.

  Throttling only attenuates the problems, but does not solve the problems. Use throttling in contention with the QoS feature otherwise it can impact service level agreement (SLA) terms.

Several practical measures that can help avoid resource starvation include:

- Designing application with QoS and throttling built in

  The service provider should design its applications with a QoS and request throttling built in. For example, if a service provider provides a facility for OS provisioning, the rate of fulfillment of the request should be based on the system current workload. Optionally, certain clients can have a higher fulfillment priority.

- Spreading out background jobs

  In multitenant architecture, housekeeping jobs usually are done on thousands of system, instead of just a few systems. Making sure that background jobs are spread across a wide spectrum of time is very important to avoid unnecessary peak loads. For example, if backup service is offered to tenants, the start of the backup job should be spread over a time period, instead of starting at the same time for all tenants.

- Detailed monitoring

  The most critical aspect on maintaining good performance is to have good infrastructure monitoring. The monitoring provides information where more resources are necessary, and is the best sizing tool to determine when more capacity is necessary.

3

Provisioning

In this chapter, we provide the basic concepts of provisioning and considerations in the multitenant AIX environment. We describe the tools for automation and implementation examples with these tools.

This chapter contains the following topics:

3.1 Introduction to provisioning

In this section we discuss how to implement provisioning in a multitenant computing environment on AIX and Power Systems.

We define resources as being hardware and software that can be shared by users, including server, processor, memory, I/O, network, storage, OS image, application software image, and so on.

The provisioning process consists of the five steps listed Table 3-1.

Table 3-1 Provisioning process steps

Step name	Flow of resources
3.3.1, "Planning" on page 63	Specify the resources we can provide
3.3.2, "Allocation" on page 64	Allocate the resources
3.3.3, "Installation" on page 66	Put together allocated resources
3.3.4, "Customization" on page 67	Customize the resources if necessary
3.3.5, "Deprovisioning" on page 69	Recycle the resources

In a multitenant environment, all steps require sophisticated management. Although we can implement each step manually, a multitenant environment has a huge number of resources and requirements. Automation is essential in the management of resources to meet the various requirements. Automation minimizes the complexity of resources and risk of management.

Note: This book covers only server provisioning. The resource pool has servers, storage, applications and various network devices such as switches, hubs and VLANs. Provisioning of storage, applications and network devices is out of the scope of this book.

For more information about provisioning and automation concepts in IBM, read about how to provision system resources according to business demand in *On Demand Operating Environment: Managing the Infrastructure,* SG24-6634.

3.2 Conceptual provisioning model

As a service provider, you are providing a multitenant utility computing service for the tenant. The tenant is the service provider's customer. Consider the following tasks to implement provisioning:

1. Ensure that enough resources are available o meet your tenant's request.
2. Perform resource planning, resource allocation, installation, customization, and deprovisioning
3. Rebalance resources, if necessary, to free up spare capacity.
4. Perform efficient job ordering.

3.2.1 Components of the conceptual model

To meet the considerations in the previous list, we can transform the provisioning steps to a conceptual model, as shown in Figure 3-1 on page 60. In this model, the flow of resources may be sequential or circular. If we have a simple resource set and simple requirement, then flow of resources is simpler. When considering resource reallocation or mobility, flow of resources is more complex.

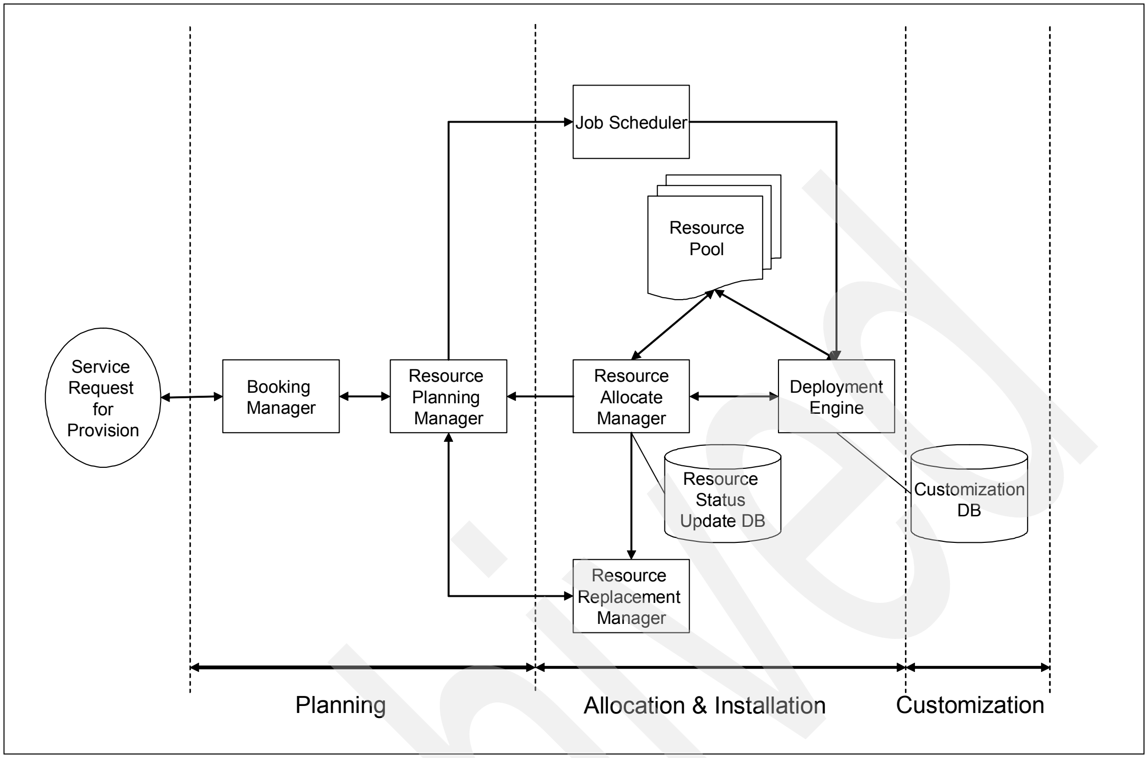
Job Scheduler
Resource Pool
Service Request for Provision
Booking Manager
Resource Planning Manager
Resource Allocate Manager
Deployment Engine
Resource Status Update DB
Customization DB
Resource Replacement Manager
Planning
Allocation & Installation
Customization

Figure 3-1 Conceptual provisioning model and steps

The conceptual model consists of several components, and each component has its own responsibility. The components and their responsibilities are described in the following paragraphs.

Booking Manager

Booking Manager is a part of the process that is responsible for the tenant request and the validity of request. If the request is not covered by the service terms or is not valid, Booking Manager denies the service request.

Resource Planning Manager

Resource Planning Manager creates the resource plan based on the information given by the Resource Allocate Manager and the Resource Replacement Manager. It decides on the viability of resource allocation, checks for enough resources and considers procuring new resource for future use.

Resource Allocation Manager

Resource Allocation Manager manages the current resource status and allocates resources. It allocates CPUs, memory, disk, network according to the resource allocation decisions made already. It makes LPARs or allocates dedicated servers. If not enough resources are available, it alarms Resource Planning Manager and Resource Replacement Manager, which makes the decision of reorganization of resources or resource mobility.

Resource Replacement Manager

If not enough resources are available in one server or resource pool, Resource Replacement Manager considers the following options:

- Moving workloads to another server or resource pool
- Adjusting resources assigned to existing workloads

Resource Replacement Manager decides resource replacement based on the information by Resource Allocate Manager.

Resource Planning Manager, Resource Allocation Manager, and Resource Replacement Manager cooperate to meet the business goals and the service level agreement (SLA).

Job Scheduler

After resource planning finishes, the Resource Planning Manager sends the deployment request to the Job Scheduler. The Job Scheduler manages the job queue for deployment.

Deployment Engine

The Deployment Engine installs the OS image and configures a tenant's resources as per the customization policy. It gets specific information from a customization database to meet various requests. After all deployment finishes, it removes the installation connection (such as network or storage) from Deployment Engine. After this step, the tenant can have an isolated OS image for the tenant's own use.

3.2.2 Automation tools

In a multitenant environment, resource management is very complex. To reduce costs and make full use of resources, automation tools are essential.

Many tools and combinations of tools exist for provisioning. They can be used in different phases of the provisioning process and their automation levels are different. For example, we can have one tool to create a new machine and another tool to install an OS.

In Figure 3-2, we map the automation software to each component. The specific implementation is discussed in 3.4, “Planning tools” on page 69, 3.5, “Allocation tools” on page 75, 3.6, “Installation tools” on page 111, and 3.7, “Customization tools” on page 130.

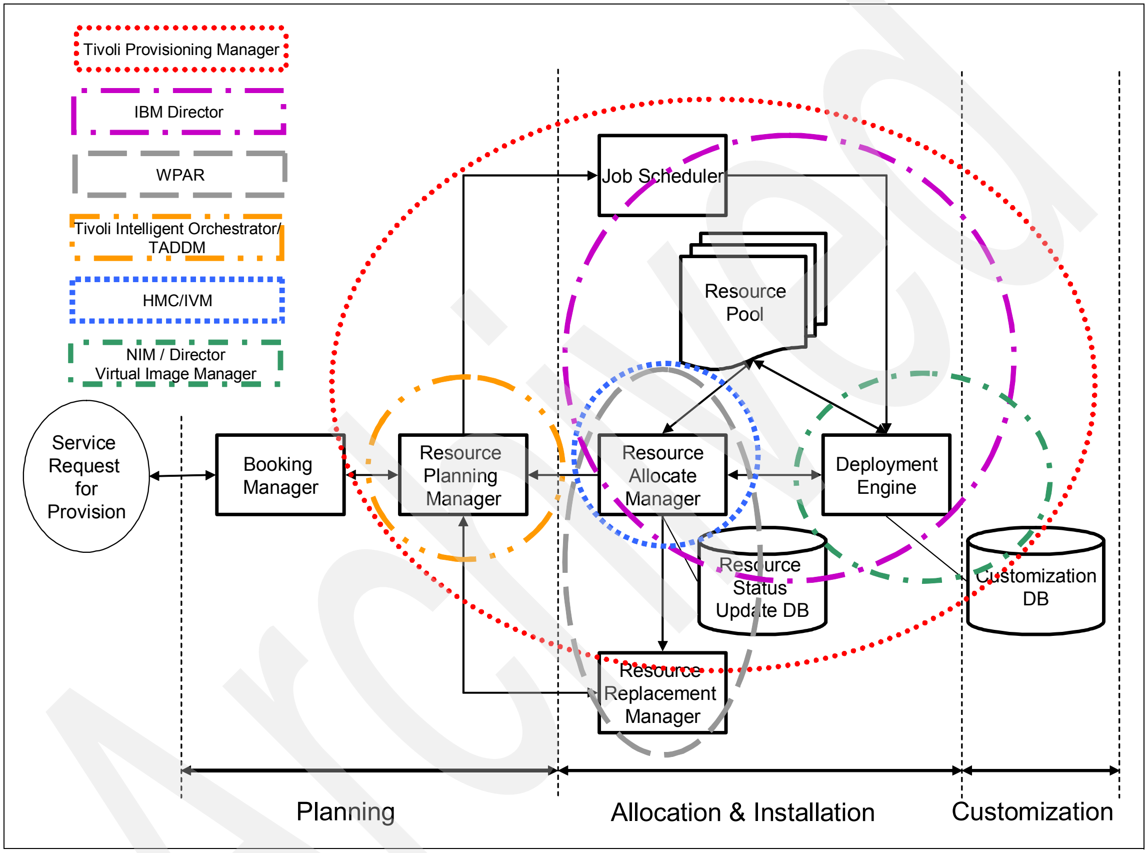

Tivoli Provisioning Manager
IBM Director
WPAR
Tivoli Intelligent Orchestrator/ TADDM
HMC/IVM
NIM / Director Virtual Image Manager
Service Request for Provision
Booking Manager
Resource Planning Manager
Job Scheduler
Resource Pool
Resource Allocate Manager
Deployment Engine
Resource Status Update DB
Customization DB
Resource Replacement Manager
Planning
Allocation & Installation
Customization

Figure 3-2 Mapping to automation tool

3.3 Considerations for provisioning

We have five steps for provisioning, as listed in Table 3-1 on page 58. They are planning, allocation, installation, customization, and deprovisioning. In this section, we provide considerations for each step when we implement a multitenant environment.

3.3.1 Planning

In the planning step, we must set our goal of service level exactly. To meet this objective, we consider the resource plan previously implemented on that system.

In a multitenant environment, we have to manage a huge number of platforms as resources. We must know the features of each resource. With this information, we decide the service level that we can provide. We also get irregular requests of resources from multiple sources. Sources can be tenants or another modules in our provisioning configuration.

For each service provided in the *service catalog* of the SLA, the following considerations are common for resource planning in the multitenant environment:

- How to manage resources efficiently
- How to reduce management cost
- How to provide for tenant requests within specified time frames (for each SLA)
- How to make use of automation tools
- How to manage the peak workload or over commitment of resources, such as servers, storage, network bandwidth, and electricity
- How to provide the tenants a secured environment from each other's environment
- How to prepare the resource capacity to support continuous service

The initial planning provides a starting point for the resource requirements. While a tenant is using their servers after installation of the OS, other tenants might demand more resources or might require fewer resources to save costs. Update the plans to meet the tenant's requests. Sometimes a server has downtime for an unplanned upgrade, and we may have to consider partition mobility. We also have to use servers that are as energy efficient as possible.

IBM Tivoli automation tools that can help with resource planning are:

- Tivoli Provisioning Manager (TPM) provides a framework for automated provisioning, configuration and maintenance.

  http://www.ibm.com/software/tivoli/products/prov-mgr/

- Tivoli Intelligent Orchestrator (TIO) provides a service level monitoring and feedback mechanism that triggers and regulates TPM activities

  http://www.ibm.com/software/tivoli/products/intell-orch/

- Tivoli Application Dependency Discovery Manager (TADDM) provides a framework for discovery of application dependencies

  http://www.ibm.com/software/tivoli/products/taddm/

With these tools, we set the goal of service level or reset the goal according to a changing environment. These tools may be combined with allocation or installation tools.

3.3.2 Allocation

In a multitenant environment, we have to automate the allocation of resources because a large number of servers and resources leads to complicated management.

Consider the following issues regarding allocation:

- How many platforms we can support
- How to allocate proper resources
- How to automate the allocation

To reduce costs and to manage resources efficiently, we can make various tools. The level of automation is very different for each tool, but the following basic features are required to support allocation for a multitenant environment:

- Create a server or logical partition (LPAR)
- Destroy a server or LPAR
- Add processors, memory, or I/O
- Remove processors, memory, or I/O
- Move processors, memory, or I/O

If we allocate resources for a tenant, we must consider how to best make use of all the resources. If our environment cannot support virtualization, we have to buy all the physical resources such as network adapters, hard disks, and I/O drawers for each tenant.

We must also consider the resource granularity our environment can support. For example, a tenant might require a very isolated and high-performance environment for their database server. In this case, we have to provide a dedicated network and high-performance server instead of a virtualized environment. To support various services, the environment must provide high granularity of resources.

The following automation tools can help with resource allocation:

- Hardware Management Console (HMC)
- Integrated Virtualization Manager (IVM)
- Tivoli Provisioning Manager (TPM)
- Workload Partition (WPAR)
- IBM Director

Use the tools to allocate various resources. Select any tool but consider the service level that you want to provide.

Resource allocation options

The resource allocation step has many options to provide to tenants. Consider the following questions regarding the tenants:

- What level of isolation is necessary?
- Is a high-availability environment required?
  - How long a break of service is acceptable during partition mobility?
  - Does the tenant want High-Availability Cluster Multi-Processing (HACMP) to be provided?
- What kind of workload does the tenant have?
  - Is the workload a processor-intensive job?
  - Is the workload memory-intensive job?
  - Does the workload require high bandwidth of network I/O?
  - Does the workload require high bandwidth of disk I/O?
  - Do the tenant want any special device for a job?

The options are also shown in Figure 3-3 on page 66.

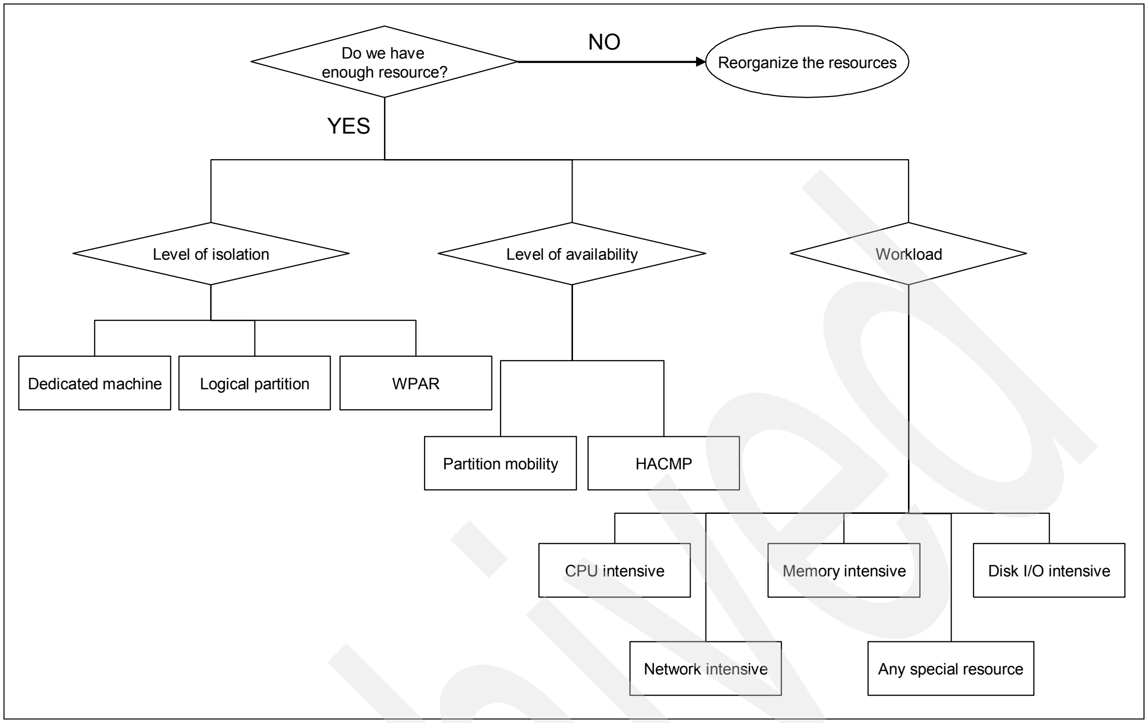

Figure 3-3 Resource allocation options for tenant

The starting point of making a decision is whether we have enough resources to meet tenant's request. If the answer is no, we can reorganize the resources we have to provide service to tenant. This decision is based on the Resource Status Update database, introduced in Figure 3-1 on page 60.

These options can be updated according to changing environment of the service provider. If the tenant does not know the options well, we can provide a default resource allocation option.

3.3.3 Installation

After resource allocation, we have to install the operating system (OS). In a multitenant environment, we have to consider the system configuration that is supporting the OS installation. Generally, OS installation is time-consuming, and requires highly efficient network and storage. Sometimes it requires distributed system configuration to provide less burden in the network.

Note: System architecture for the installation process depends on how many tenants we have. We can have single installation server or multiple installation servers with hierarchical structure. We do not cover all the different scenarios in this book.

Many issues can be considered in OS installation except system configuration:

- How to make a base OS installation image
- How many OS images we can provide
- How to manage OS image versions
- How can we automate the installation

After OS installation, we have to customize each environment. If our base OS image has few common features installed, customization will take a long time to install additional packages. If our OS image has many common features, customization can be difficult while removing unnecessary packages for a tenant.

To reduce costs and to manage resources efficiently, we can use various tools. The level of automation is very different for each tool. Basic tool features for installation of a multitenant environment should include:

- Install OS
- Start OS
- Shutdown OS
- Retain OS images for future use

Automation tools that can help with installation and management of the OS are:

- Network Installation Management (NIM)
- Tivoli Provisioning Manager (TPM)
- IBM Director Virtual Image Manager

3.3.4 Customization

In a multitenant environment, there are various requests for configuration changes by tenants. There is also the trade-off about the degree of customization between predefined services and personalized service. Customization considerations include:

- What level of customization to provide
- How to automate the customization

After the installation, the tenant has an IP address, a user name and a password to be able to log in to the system and start using the new equipment. For this reason, certain customizable information has to be executed by the service provider to deliver a unique environment for the client.

For example, several major customization tasks and their properties are:

- Network
  - Configure general IP parameters
    - IP addresses
    - Netmasks
    - Host name and default gateway IP address
    - Static routes if needed
  - Import or add entries to an existing `/etc/hosts` file
- Users and groups
  - Create new users
  - Create new groups
- Environment configurations
  - General date information, such as date, time, and time zone
- User environment and files
  - Setting of particular language, environment variables (`/etc/profile` or `/etc/environment`)
  - Setting other customizable files (`/etc/motd`)
- Operating system technology levels and bundles
  - Installation of technology levels and service packs
  - Installation of new software bundles
- Management product installation
  - Accounting agent
  - Monitoring agent
  - Security agent
  - Managing agent
- User application installation
  - Installation of particular applications and tools

Note: The customized application installation is covered briefly in this book.

Automation tools that can automate customization of various properties are:

- Network Installation Management (NIM)
- Tivoli Provisioning Manager (TPM)

After customization, the service provider can use different ways to inform the tenant of the services. The information can be items such as user ID, password, IP address, simple FAQ, and how to connect the server. This information can be communicated to the tenant by any of the following methods:

- E-mail
- Web
- SMS
- Fax
- Letter

3.3.5 Deprovisioning

In the resource allocation step, you must consider deprovisioning plan for resource recycling. Considerations for deprovisioning are:

- Retain accounting data for billing
- Retain performance data for capacity planning and customization

Basic features to support customization for multitenant environment are:

- Remove files and file systems
- Scrub used storage
- Unconfigure server or LPAR

> **Note:** Although this book does not cover deprovisioning steps, deprovisioning can be implemented by automation tools.

3.4 Planning tools

We have several automation tools that can support the planning step. In this section, we discuss the various tools and the usage.

3.4.1 Planning with Tivoli Provisioning Manager

Tivoli Provisioning Manager (TPM) can discover the existing system configuration information. We can develop resource planning based on the information. TPM can manage the various resources, discover changes, and automatically update the resource information.

TPM provides **Software Management** and **Inventory** menus for planning, as shown in Figure 3-4. With this menu, we can see our hardware and software status information.

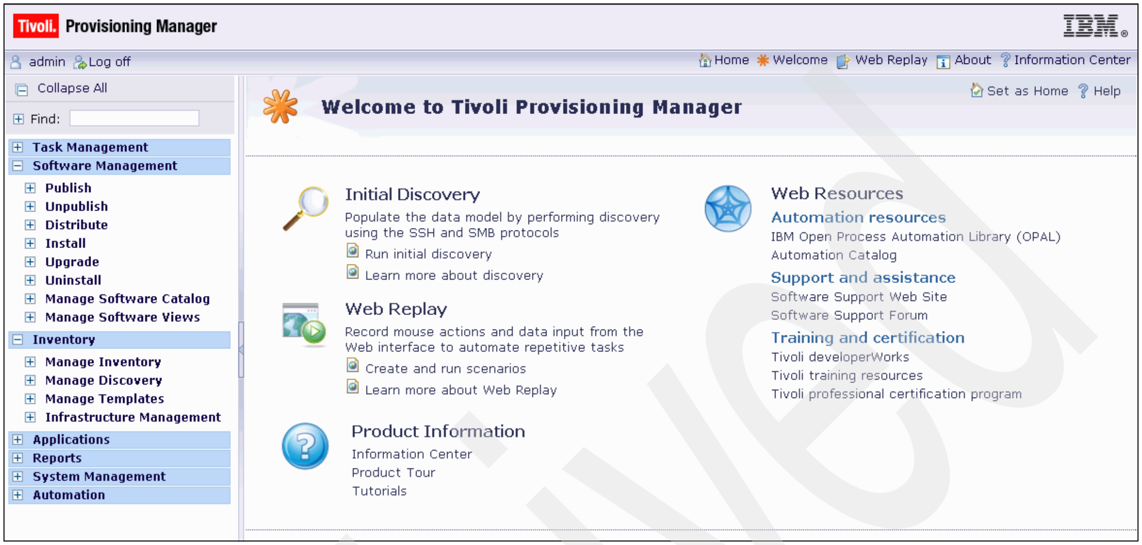

Tivoli. Provisioning Manager

admin Log off

Home Welcome Web Replay About Information Center

Collapse All

Find:

- Task Management
- Software Management
  - Publish
  - Unpublish
  - Distribute
  - Install
  - Upgrade
  - Uninstall
  - Manage Software Catalog
  - Manage Software Views
- Inventory
  - Manage Inventory
  - Manage Discovery
  - Manage Templates
  - Infrastructure Management
- Applications
- Reports
- System Management
- Automation

Set as Home Help

Welcome to Tivoli Provisioning Manager

Initial Discovery

Populate the data model by performing discovery using the SSH and SMB protocols

Run initial discovery

Learn more about discovery

Web Replay

Record mouse actions and data input from the Web interface to automate repetitive tasks

Create and run scenarios

Learn more about Web Replay

Product Information

Information Center

Product Tour

Tutorials

Web Resources

Automation resources

IBM Open Process Automation Library (OPAL)

Automation Catalog

Support and assistance

Software Support Web Site

Software Support Forum

Training and certification

Tivoli developerWorks

Tivoli training resources

Tivoli professional certification program

Figure 3-4 Menus provided by TPM

If you have a new Power Systems server, use either of the following methods to discover that system:

- Select **Inventory** → **Manage Discovery** → **Initial Discovery**.
- Use **Hardware Management Console Discovery**.

  See 3.5.3, "Allocation with Tivoli Provisioning Manager" on page 80, for information about the Hardware Management Console Discovery.

After discovering the new system, manage the status of the system by selecting **Inventory** → **Manage Inventory** as shown in Figure 3-5 on page 71.

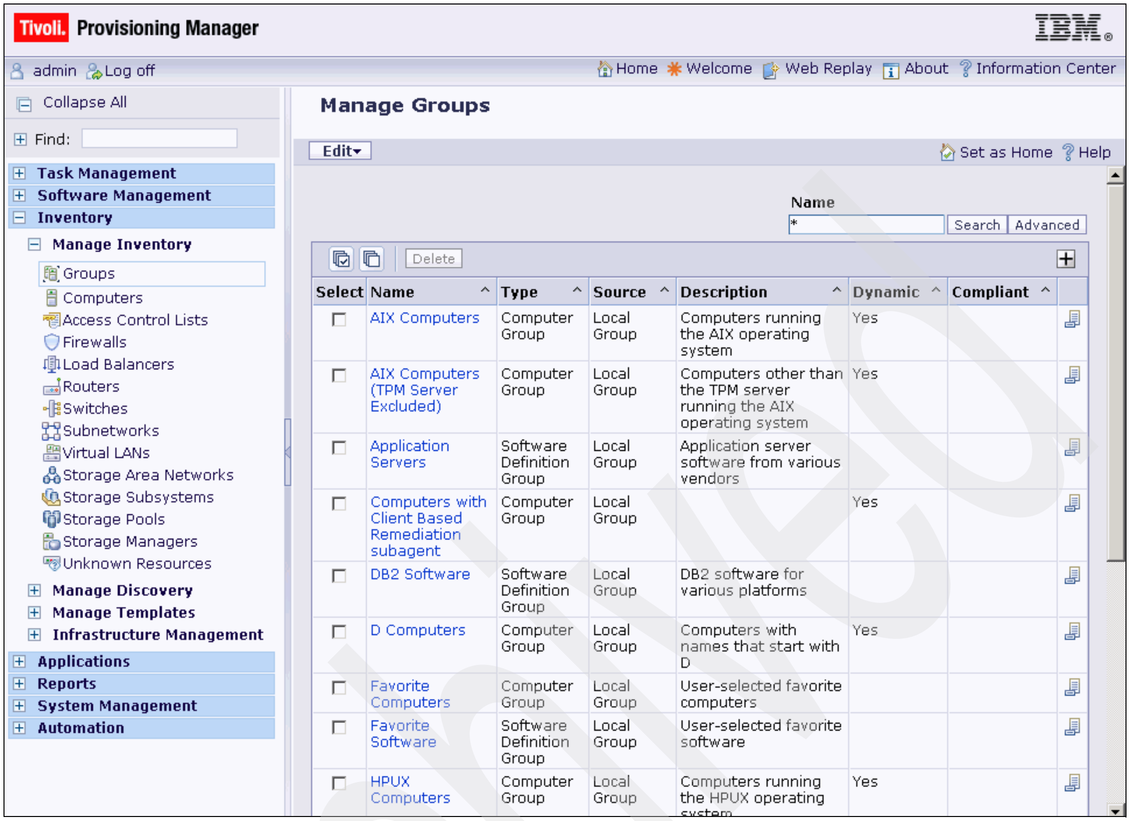

- Task Management
- Software Management
- Inventory
  - Manage Inventory
    - Groups
    - Computers
    - Access Control Lists
    - Firewalls
    - Load Balancers
    - Routers
    - Switches
    - Subnetworks
    - Virtual LANs
    - Storage Area Networks
    - Storage Subsystems
    - Storage Pools
    - Storage Managers
    - Unknown Resources
  - Manage Discovery
  - Manage Templates
  - Infrastructure Management
- Applications
- Reports
- System Management
- Automation

Manage Groups

Name: *

Select	Name	Type	Source	Description	Dynamic	Compliant
□	AIX Computers	Computer Group	Local Group	Computers running the AIX operating system	Yes	
□	AIX Computers (TPM Server Excluded)	Computer Group	Local Group	Computers other than the TPM server running the AIX operating system	Yes	
□	Application Servers	Software Definition Group	Local Group	Application server software from various vendors		
□	Computers with Client Based Remediation subagent	Computer Group	Local Group		Yes	
□	DB2 Software	Software Definition Group	Local Group	DB2 software for various platforms		
□	D Computers	Computer Group	Local Group	Computers with names that start with D	Yes	
□	Favorite Computers	Computer Group	Local Group	User-selected favorite computers		
□	Favorite Software	Software Definition Group	Local Group	User-selected favorite software		
□	HPUX Computers	Computer Group	Local Group	Computers running the HPUX operating system	Yes	

Figure 3-5 Inventory management with TPM

Each time the service provider gets new resources or any changes happen, select **Inventory** → **Manage Discovery** → **Discovery Configuration** as shown in Figure 3-6 on page 72. We can discover new resources or manage change in resources.

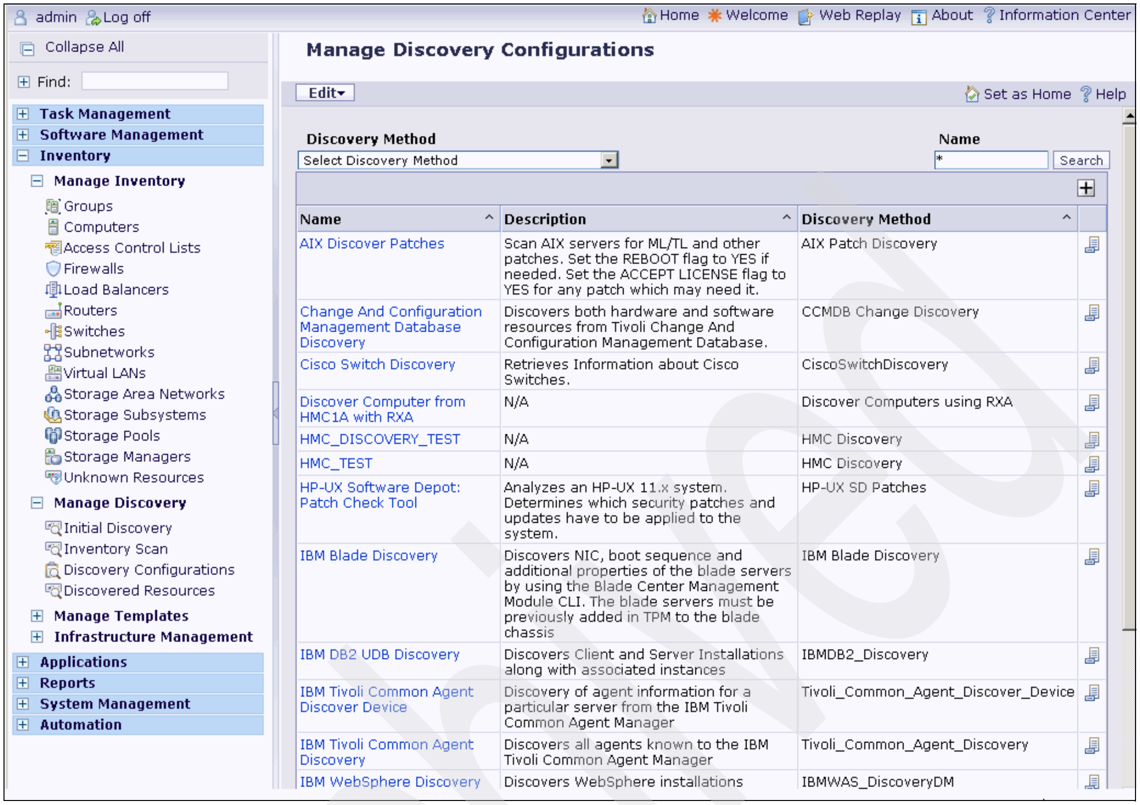

Manage Discovery Configurations

Discovery Method: Select Discovery Method

Name: *

Name	Description	Discovery Method
AIX Discover Patches	Scan AIX servers for ML/TL and other patches. Set the REBOOT flag to YES if needed. Set the ACCEPT LICENSE flag to YES for any patch which may need it.	AIX Patch Discovery
Change And Configuration Management Database Discovery	Discovers both hardware and software resources from Tivoli Change And Configuration Management Database.	CCMDB Change Discovery
Cisco Switch Discovery	Retrieves Information about Cisco Switches.	CiscoSwitchDiscovery
Discover Computer from HMC1A with RXA	N/A	Discover Computers using RXA
HMC_DISCOVERY_TEST	N/A	HMC Discovery
HMC_TEST	N/A	HMC Discovery
HP-UX Software Depot: Patch Check Tool	Analyzes an HP-UX 11.x system. Determines which security patches and updates have to be applied to the system.	HP-UX SD Patches
IBM Blade Discovery	Discovers NIC, boot sequence and additional properties of the blade servers by using the Blade Center Management Module CLI. The blade servers must be previously added in TPM to the blade chassis	IBM Blade Discovery
IBM DB2 UDB Discovery	Discovers Client and Server Installations along with associated instances	IBMDB2_Discovery
IBM Tivoli Common Agent Discover Device	Discovery of agent information for a particular server from the IBM Tivoli Common Agent Manager	Tivoli_Common_Agent_Discover_Device
IBM Tivoli Common Agent Discovery	Discovers all agents known to the IBM Tivoli Common Agent Manager	Tivoli_Common_Agent_Discovery
IBM WebSphere Discovery	Discovers WebSphere installations	IBMWAS_DiscoveryDM

Figure 3-6 Set discovery option in initial discovery

Every discovery function provides a time schedule. We can run the discovery on specific time or repeat the specified discovery, as shown in Figure 3-7 on page 73. TPM provides various notifications to the service provider.

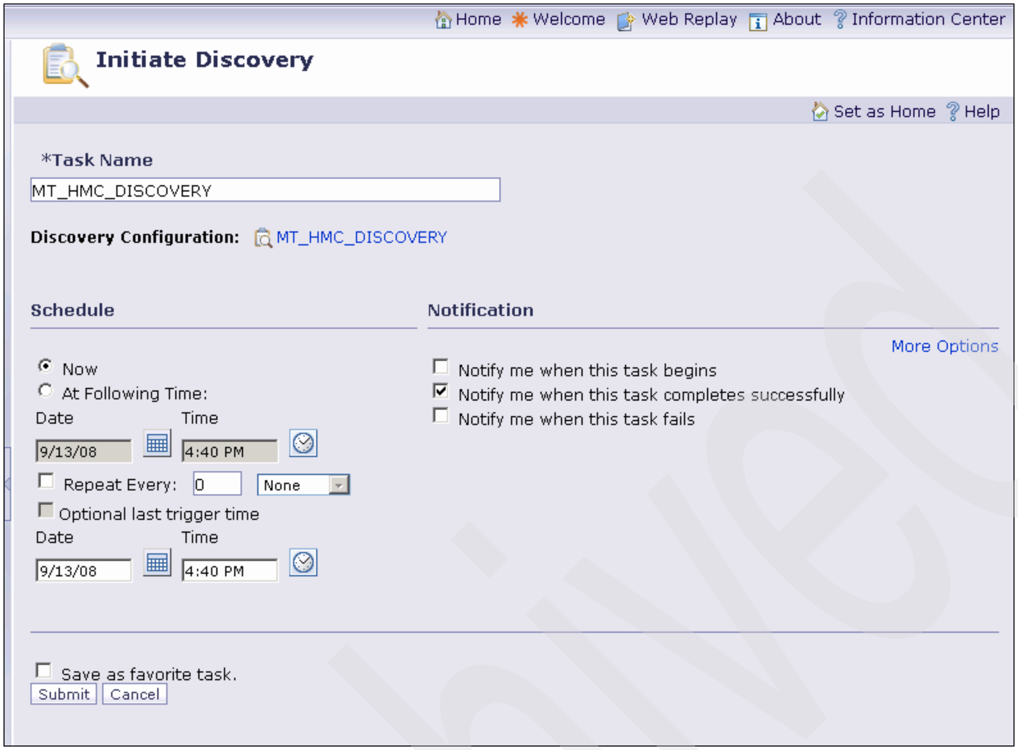

Initiate Discovery

*Task Name: MT_HMC_DISCOVERY

Discovery Configuration: MT_HMC_DISCOVERY

Schedule

● Now
○ At Following Time:

Date	Time
9/13/08	4:40 PM

□ Repeat Every: 0 None

□ Optional last trigger time

Date	Time
9/13/08	4:40 PM

Notification

□ Notify me when this task begins
☑ Notify me when this task completes successfully
□ Notify me when this task fails

□ Save as favorite task.

Figure 3-7 Various discovery option

3.4.2 Planning with Tivoli Intelligent Orchestrator

Tivoli Intelligent Orchestrator (TIO) supports policy-based orchestration that is based on TPM. Policy-based orchestration enables a service provider to optimize resources by actively sensing changes in the environment to meet business goals by predefined policy.

Based on predefined policy and priorities, TIO makes decisions about what action it has to take, when to start deployment, and which resources to use. Internally, TIO operates in a loop. Phases in the loop are:

Monitor TIO monitors various metrics within an system environment and compares them to thresholds. If the thresholds are exceeded, TIO can send events to external systems, to the Analyze phase, or both.

Analyze After receiving events from the Monitor phase, TIO correlates the events as necessary, analyses the cause of the notification, and sends the results to the Planning phase, if necessary.

Plan Based on the results of the Analyze phase, TIO finds out a solution, such as rolling back a software upgrade, adding another server, or adding more storage space to a database.

Execute Implements the solution received from the Plan phase. TIO invokes TPM to provide this execution.

You may choose one of these modes for each phase:

- Manual mode

  This mode requires an operator to manually start workflows to take action.

- Semiautomatic mode

  This mode queries the operator for permission to execute a set of workflows based on a predefined orchestrated events.

- Automatic mode

  This mode executes the orchestration action automatically, requiring no operator intervention.

For more information, see *Automated provisioning of resources for data center environments*, G507-1090 and *Provisioning On Demand Intoroducing IBM Tivoli Intelligent Think Dynamic Orchestrator,* SG24-8888.

3.4.3 Planning with Tivoli Application Dependency Discovery Manager

Tivoli Application Dependency Discovery Manager (TADDM) provides complete visibility for complexity of business applications by automatically creating and maintaining application infrastructure maps. The comprehensive TADDM application maps can display automatically discovered IT resources and can include complete business application dependencies, detailed configuration values, and an accurate change history graphically.

TADDM creates a topological definition of applications, which can be used by other management applications like TPM. If you provide complex software structure for a tenant, this map is a starting point of resource planning.

See *Deployment Guide Series: IBM Tivoli Application Dependency Discovery Manager V7.1,* SG24-7616.

3.5 Allocation tools

Several tools can support the allocation step. Some tools provide more automation than others. This section discusses the various tools and their usage.

3.5.1 Allocation with Hardware Management Console

The Hardware Management Console (HMC) is used to control the Power Systems servers. It uses a separated, private network to manage the servers and can be accessed through a Web browser. Figure 3-8 shows an HMC connected to two Power Systems servers.

Use the HMC to:

- Create or delete LPARs
- Dynamically add, move or remove CPU, memory, physical and virtual devices
- Start up and shut down the managed system
- Start up and shut down the OS that is within a particular LPAR

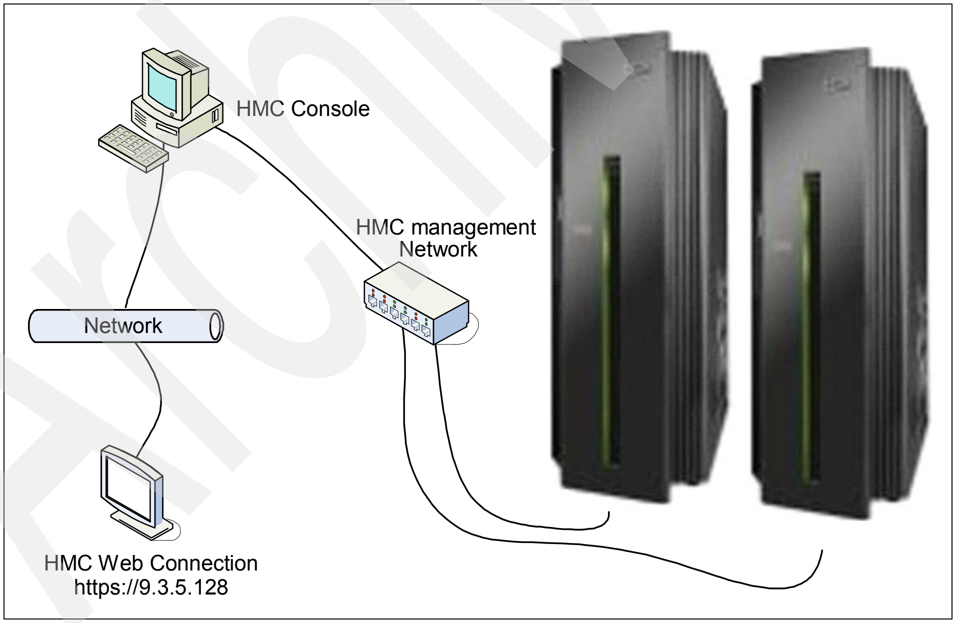

Figure 3-8 An HMC connection with two Power Systems servers

To access the HMC's Web interface, use the following format to enter a Web address in the browser:

```
https://<HMC IP or Hostname>
```

A window from an HMC Web interface as shown in Figure 3-9.

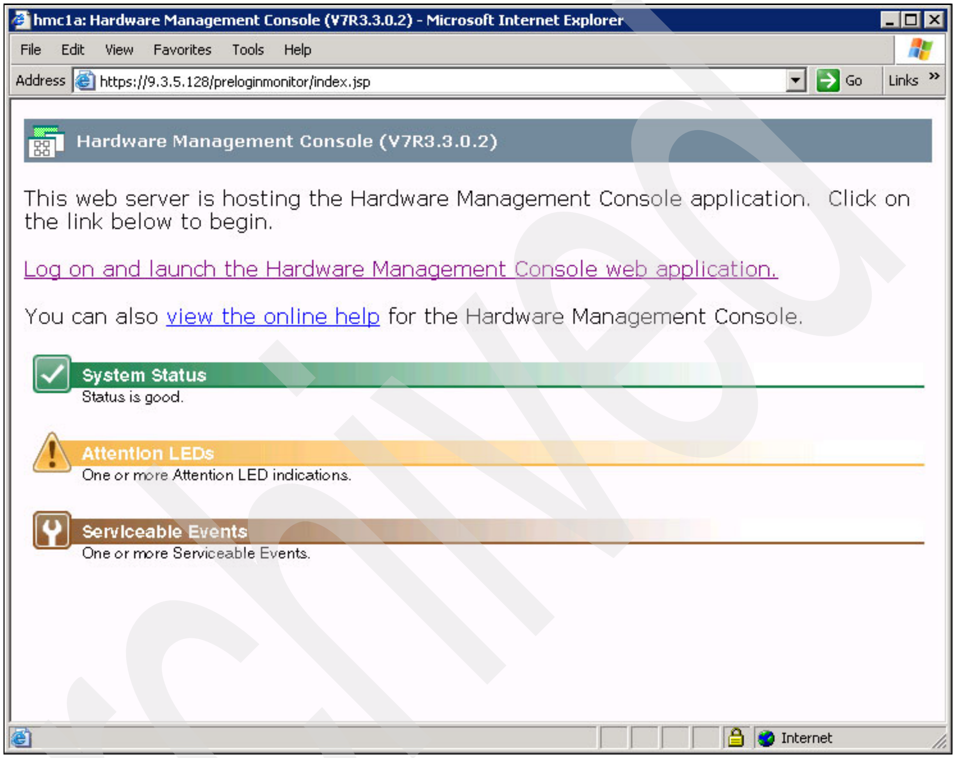
hmc1a: Hardware Management Console (V7R3.3.0.2) - Microsoft Internet Explorer
File
Edit
View
Favorites
Tools
Help
Address
https://9.3.5.128/preloginmonitor/index.jsp
Go
Links

Hardware Management Console (V7R3.3.0.2)

This web server is hosting the Hardware Management Console application. Click on the link below to begin.

Log on and launch the Hardware Management Console web application.

You can also view the online help for the Hardware Management Console.

System Status
Status is good.

Attention LEDs
One or more Attention LED indications.

Serviceable Events
One or more Serviceable Events.

Figure 3-9 Hardware Management Console from HMC V7

After logging on, select the managed server. A window is displayed, as shown in Figure 3-10 on page 77.

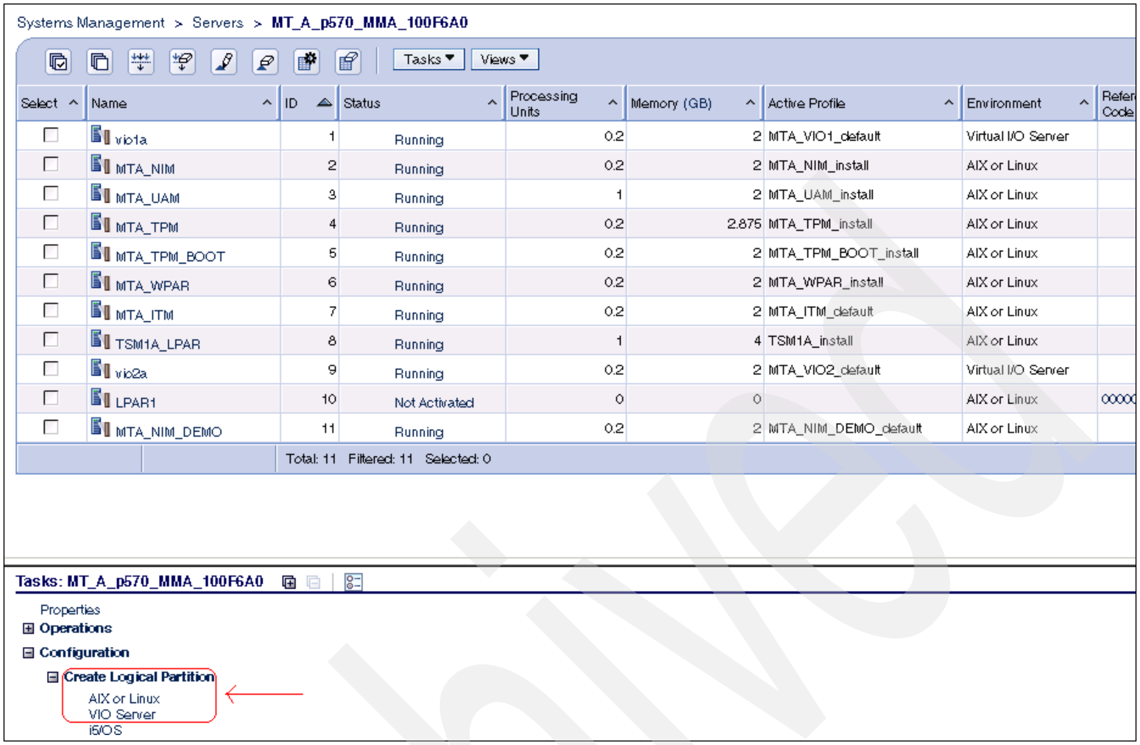

Systems Management > Servers > MT_A_p570_MMA_100F6A0

Select	Name	ID	Status	Processing Units	Memory (GB)	Active Profile	Environment	Refer Code
□	vio1a	1	Running	0.2	2	MTA_VIO1_default	Virtual I/O Server	
□	MTA_NIM	2	Running	0.2	2	MTA_NIM_install	AIX or Linux	
□	MTA_UAM	3	Running	1	2	MTA_UAM_install	AIX or Linux	
□	MTA_TPM	4	Running	0.2	2.875	MTA_TPM_install	AIX or Linux	
□	MTA_TPM_BOOT	5	Running	0.2	2	MTA_TPM_BOOT_install	AIX or Linux	
□	MTA_WPAR	6	Running	0.2	2	MTA_WPAR_install	AIX or Linux	
□	MTA_ITM	7	Running	0.2	2	MTA_ITM_default	AIX or Linux	
□	TSM1A_LPAR	8	Running	1	4	TSM1A_install	AIX or Linux	
□	vio2a	9	Running	0.2	2	MTA_VIO2_default	Virtual I/O Server	
□	LPAR1	10	Not Activated	0	0		AIX or Linux	0000
□	MTA_NIM_DEMO	11	Running	0.2	2	MTA_NIM_DEMO_default	AIX or Linux	

Total: 11 Filtered: 11 Selected: 0

Tasks: MT_A_p570_MMA_100F6A0

- Properties
- ⊞ Operations
- ⊟ Configuration
  - ⊟ Create Logical Partition
    - AIX or Linux
    - VIO Server
    - i5/OS

Figure 3-10 Managed server partitions

We can create a new partition on HMC, and run AIX, Linux, or Virtual I/O Server. Select the configuration type under the **Create Logical Partition** task. The Create Partition panel is displayed, as shown in Figure 3-11 on page 78.

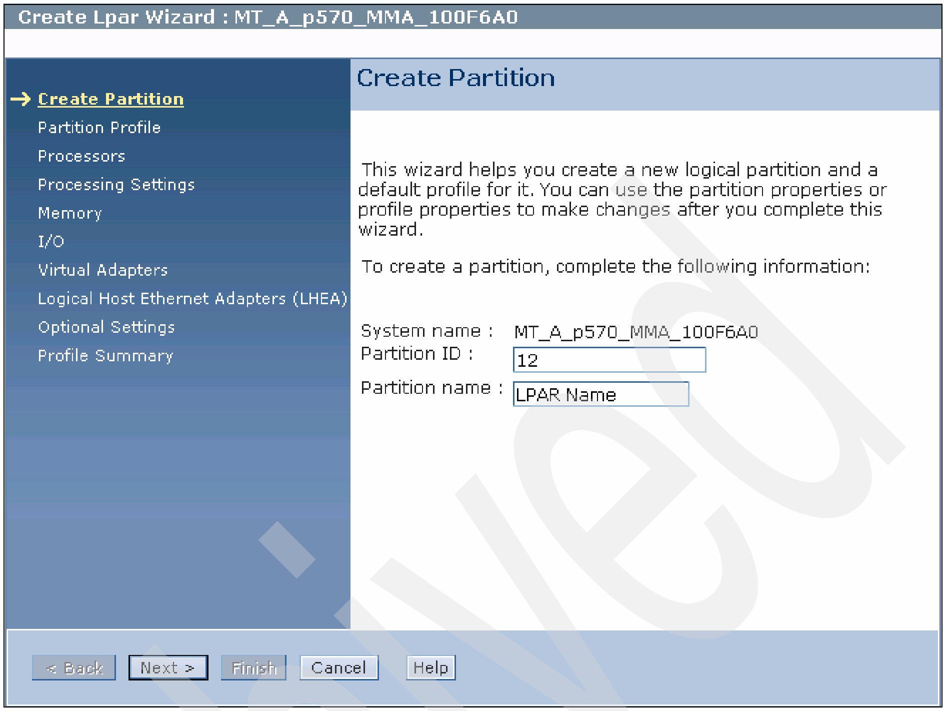

Create Partition

This wizard helps you create a new logical partition and a default profile for it. You can use the partition properties or profile properties to make changes after you complete this wizard.

To create a partition, complete the following information:

System name : MT_A_p570_MMA_100F6A0
Partition ID : 12
Partition name : LPAR Name

Figure 3-11 Create Partition panel in HMC

We fill in all the information about the LPAR that we are creating. We can delete, activate, or change the configuration of a current LPAR with other main tasks.

Basically, the HMC provides manual tasks to create and manage new machines, but can provide automation tasks in combination with the data supplied by the system planning tool (SPT).

See the following resources:

- For information about SPT:

  https://www.ibm.com/systems/support/tools/systemplanningtool/

- For details about HMC and its functions, see the IBM Systems Hardware Information Center:

  http://publib.boulder.ibm.com/infocenter/eserver/v1r3s/index.jsp

3.5.2 Allocation with Integrated Virtualization Manager

The Integrated Virtualization Manager (IVM) provides hardware management by using certain HMC features, but provides more simplified and limited functions.

IVM is one part of the Virtual I/O Server features. If you plan to use the IVM, the Virtual I/O Server software must be installed. IVM is managed by a Web browser interface and also provides a command line option. Normally, IVM is used to manage small configurations.

The IVM enables you to:

- Create or delete an LPAR
- Dynamically add or remove CPU, memory and only virtual devices
- Startup and shutdown the managed system
- Startup and shutdown individual LPARs

The IVM managed system cannot be managed by the HMC at the same time. Basically the hardware is configured not to use HMC and installed with the Virtual I/O Server media. In Figure 3-12 shows the IVM configuration.

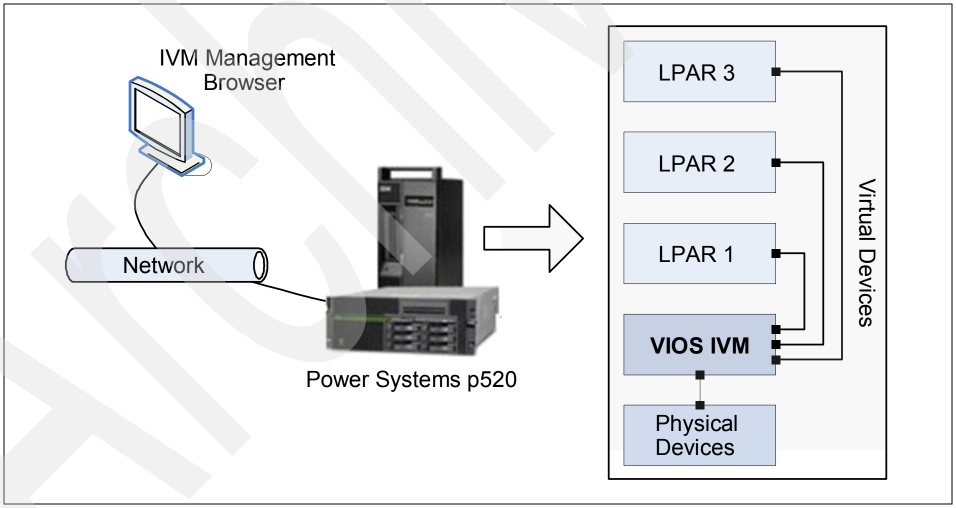

Figure 3-12 Basic IVM schema

For more details about IVM, see *PowerVM Virtualization on IBM system p: Introduction and Configuration Fourth Edition,* SG24-7940.

3.5.3 Allocation with Tivoli Provisioning Manager

Tivoli Provisioning Manager (TPM) provides graphical user interface and all the packaged tools for the service provider to automate resource allocation and installation process. TPM can also support the centralized management of other resources like storage, switches, and even other vendors' hardware or software.

TPM can support multiple different operating systems, physical servers, and virtualization technologies. The term *virtual server* is used in TPM for any type of virtualized server, such as an LPAR in Power Systems server. In Power Systems environment, TPM can automate the creation of the LPAR and allocation of the required resources.

For allocation with TPM, we have to install TPM on AIX first. Appendix A, "Software components used with TPM" on page 273 contains a list of specific software components and platform versions.

In this section, we discuss the following topics:

- "Tivoli Provisioning Manager test environment" on page 80
- "Starting up Tivoli Provisioning Manager" on page 81
- "Discovery of Hardware Management Console" on page 83
- "Creating and discovering a Virtual I/O Server" on page 88
- "Discovering the managed system information" on page 93
- "Creating a virtual server template" on page 98
- "Allocating a new virtual server" on page 101

Tivoli Provisioning Manager test environment

In the test environment, we have:

- TPM server tpm1a
- DNS server itm1a
- NIM server tbo1a
- Servers for tenants

The system configuration is shown in Figure 3-13 on page 81.

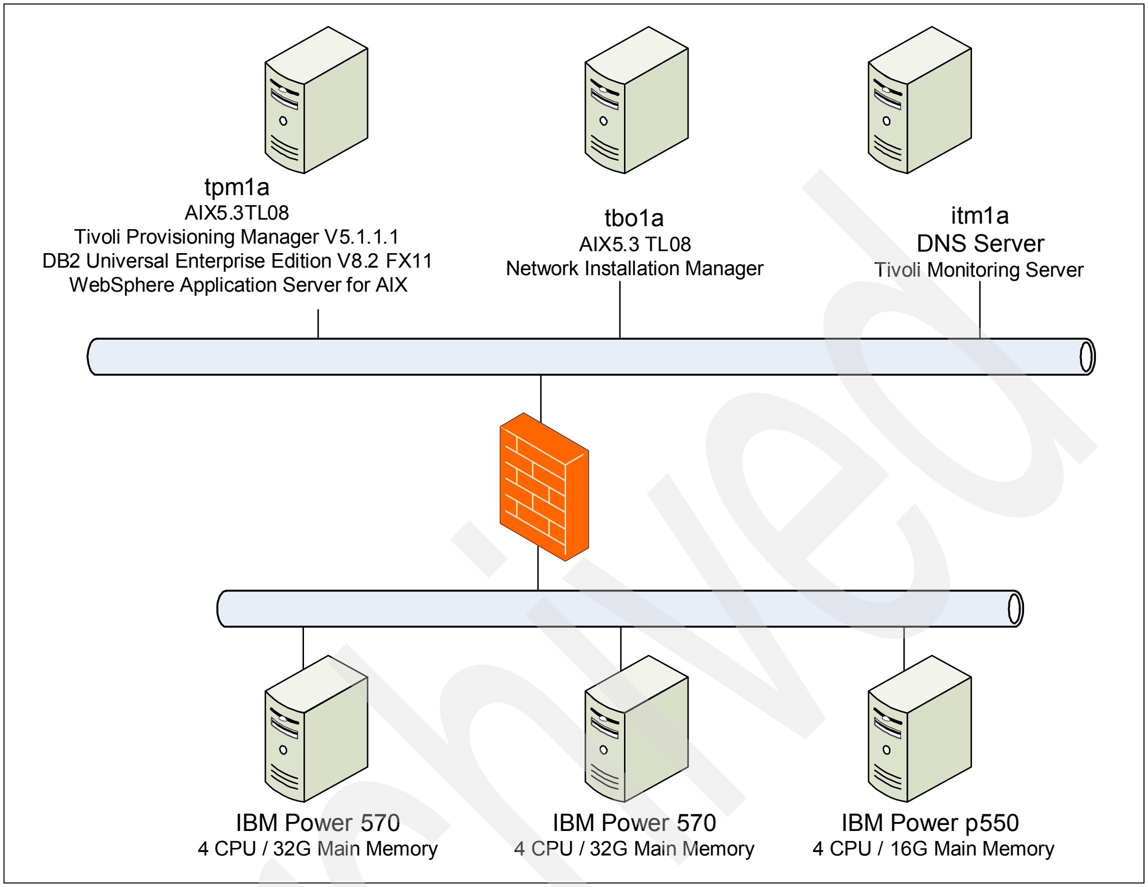

Figure 3-13 System configuration for server allocation with TPM

To configure the whole environment, TPM makes use of DNS service. To implement the automated provisioning environment, implement DNS and make the policy of IP allocation.

Starting up Tivoli Provisioning Manager

We start TPM on server with the shell command as shown in Example 3-1 on page 82.

Example 3-1 Start TPM server with shell command

```
tpm1a.austin.ibm.com:/root#>su - tioadmin
$ cd $TIO_HOME
$ cd tools
$ pwd
/opt/ibm/tivoli/tpm/tools
$ ./tio.sh start
Starting WSWB Help System...
Starting WAS on server1...

ADMU0116I: Tool information is being logged in file
           /opt/ibm/tivoli/tpm/tioprofile/logs/server1/startServer.log
ADMU0128I: Starting tool with the tioprofile profile
ADMU3100I: Reading configuration for server: server1
ADMU3200I: Server launched. Waiting for initialization status.
ADMU3000I: Server server1 open for e-business; process id is 938108

Starting Deployment Engine...
Starting Soap service...
Starting Policy Engine...
Starting Agent Shell Server...
Starting DMS Result Server...
Starting Activity Plan Engine...
TIO startup completed.
Check /var/IBM/tivoli/common/COP/logs for startup errors.
$
```

We can connect TPM with a Web browser. The format of the address is:

`https://<fully qualified name>:9045/tcWebUI`

We used the following URL:

`https://tpm1a.austin.ibm.com:9045/tcWebUI`

The user ID is `admin`, as shown in Figure 3-14 on page 83.

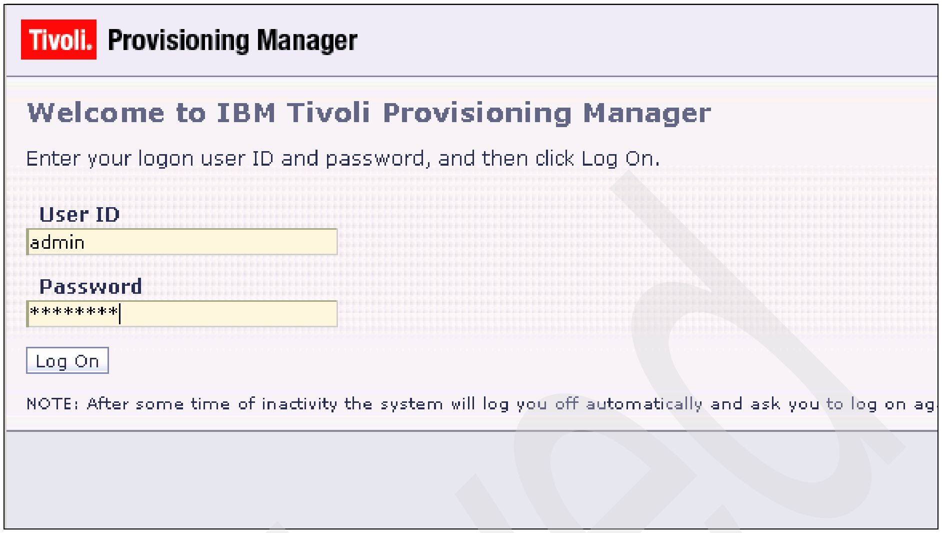

Welcome to IBM Tivoli Provisioning Manager

Enter your logon user ID and password, and then click Log On.

User ID: admin

Password: ********

NOTE: After some time of inactivity the system will log you off automatically and ask you to log on ag

Figure 3-14 TPM logon window

Discovery of Hardware Management Console

TPM must discover the HMC and host platform to manage and create the LPAR on the host platform.

As shown in Figure 3-15 on page 84, define a computer for the HMC:

1. Select **Inventory** → **Manage Inventory** → **Computers**.
2. Select **Edit** → **Add Computer**.

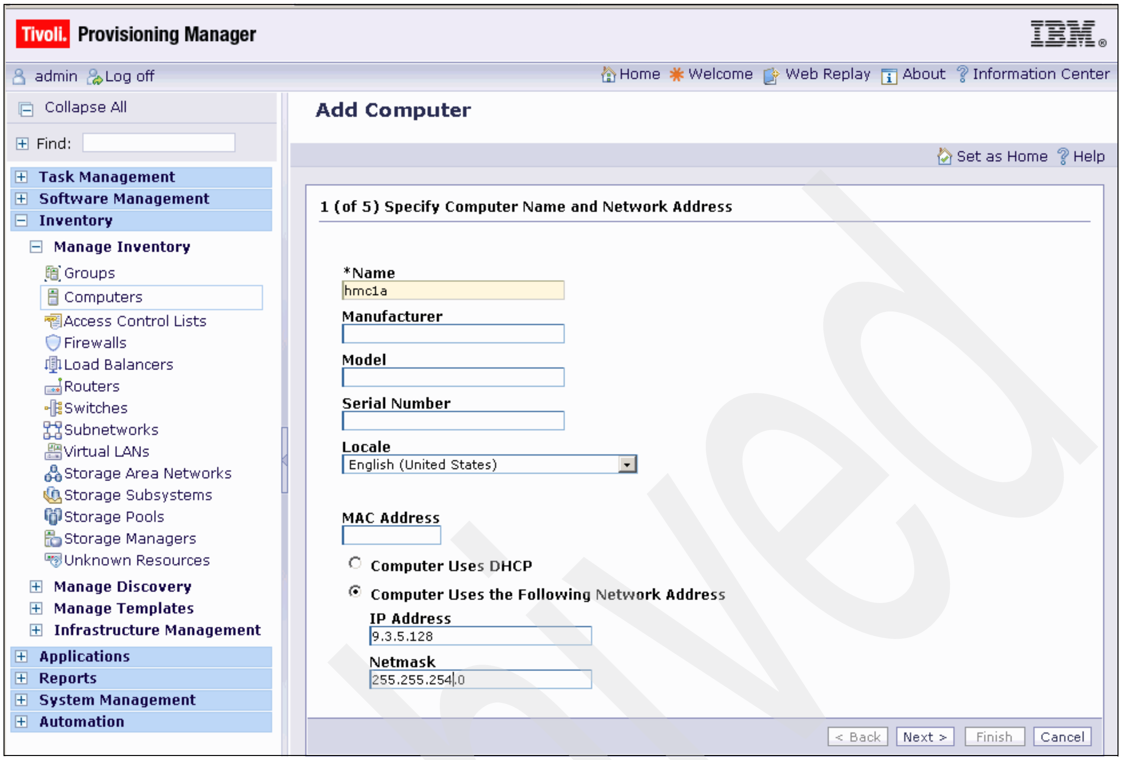

Add Computer

1 (of 5) Specify Computer Name and Network Address

*Name: hmc1a

Manufacturer:

Model:

Serial Number:

Locale: English (United States)

MAC Address:

○ Computer Uses DHCP

● Computer Uses the Following Network Address

IP Address: 9.3.5.128

Netmask: 255.255.254.0

Figure 3-15 Define a computer for the HMC in Add Computer menu

After the HMC is created:

1. Select **Edit** → **Add Credentials** to define the Service Access Point (SAP).
2. Select credential pair types as `SSH and SCP`.
3. Specify source computer as TPM server, and target computer as HMC.

For RSA credentials, enter Authentication Information. For User ID, enter `hscroot` and specify Search Key as shown in Figure 3-16 on page 85.

> **Note:** If you have another HMC, then instead of using ID as `hscroot`, use another HMC ID with the super administrator role or operator.

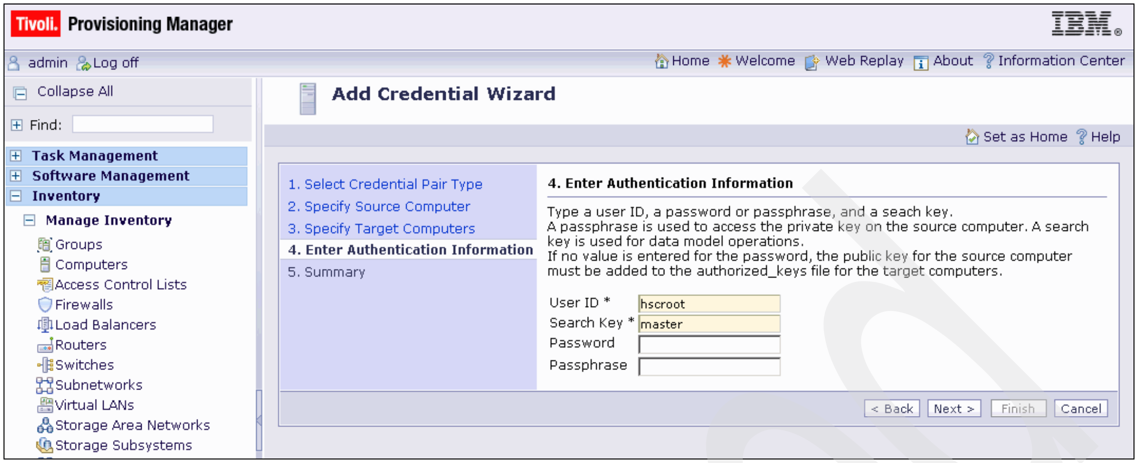

Add Credential Wizard

1. Select Credential Pair Type
2. Specify Source Computer
3. Specify Target Computers
4. Enter Authentication Information
5. Summary

4. Enter Authentication Information

Type a user ID, a password or passphrase, and a seach key.
A passphrase is used to access the private key on the source computer. A search key is used for data model operations.
If no value is entered for the password, the public key for the source computer must be added to the authorized_keys file for the target computers.

User ID * hscroot
Search Key * master
Password
Passphrase

Figure 3-16 Define RSA credentials for TPM and HMC communication

Enable the remote command on HMC. Log in to the HMC and click **HMC Management** → **Remote Command Execution**. Select **Enable remote command execution using the ssh facility** as shown in Figure 3-17.

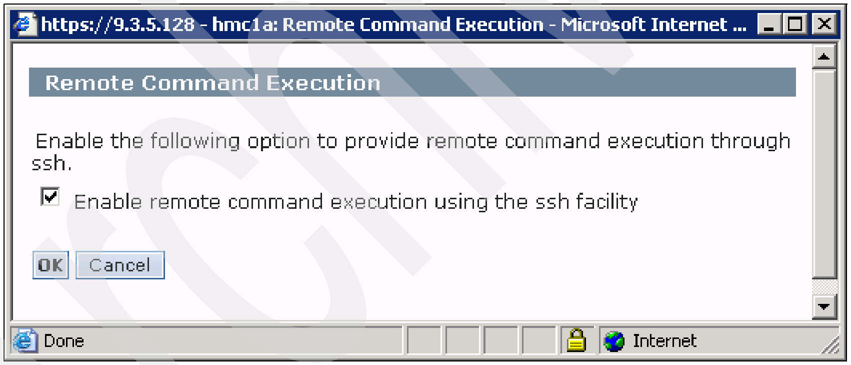
https://9.3.5.128 - hmc1a: Remote Command Execution - Microsoft Internet ...
OK
Cancel
Done
Internet

Remote Command Execution

Enable the following option to provide remote command execution through ssh.

☑ Enable remote command execution using the ssh facility

Figure 3-17 Enable Remote Command Execution in HMC

To test TPM from the HMC, run the `Device.ExecuteCommand` command.

Click **Find** and enter `Device.ExecuteCommand.` The Workflow Editor of Device.ExecuteCommand is displayed, as shown in Figure 3-18 on page 86.

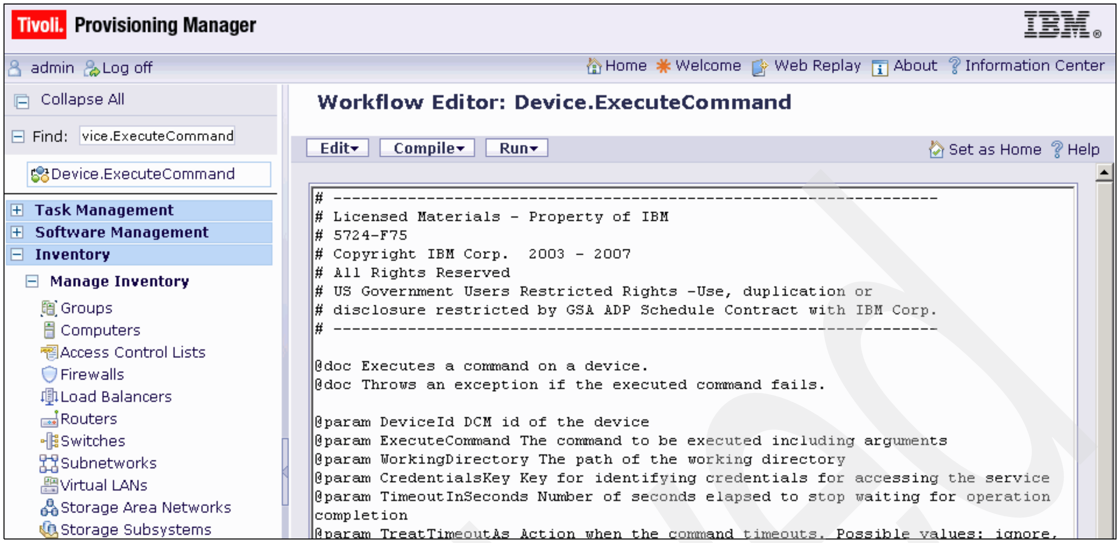

```
# -----------------------------------------------------------------
# Licensed Materials - Property of IBM
# 5724-F75
# Copyright IBM Corp.  2003 - 2007
# All Rights Reserved
# US Government Users Restricted Rights -Use, duplication or
# disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
# -----------------------------------------------------------------

@doc Executes a command on a device.
@doc Throws an exception if the executed command fails.

@param DeviceId DCM id of the device
@param ExecuteCommand The command to be executed including arguments
@param WorkingDirectory The path of the working directory
@param CredentialsKey Key for identifying credentials for accessing the service
@param TimeoutInSeconds Number of seconds elapsed to stop waiting for operation
completion
@param TreatTimeoutAs Action when the command timeouts. Possible values: ignore,
```

Figure 3-18   Workflow Editor: Device.ExecuteCommand

On the Workflow Editor page, click **Run**. Type the **device ID** for HMC computer.

To determine the Device ID, select **Inventory** → **Manage Inventory** → **Computers**, and then move the mouse over the computer name. You see the device ID. In Figure 3-19, the device ID for computer hmc1a is 11949.

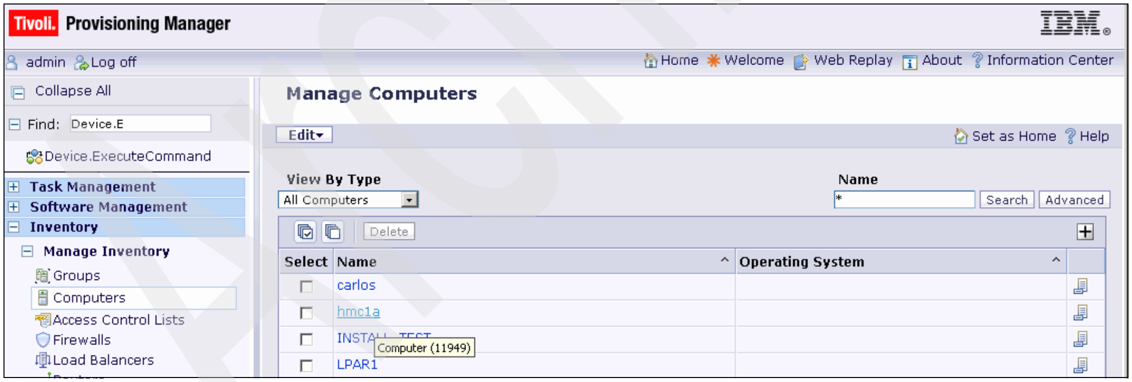

Select	Name	Operating System
□	carlos	
□	hmc1a	
□	INSTALL_TEST	
□	LPAR1	

Computer (11949)

Figure 3-19   Determine the device ID of HMC in Inventory menu

Another way to determine the device ID is to select **Inventory** → **Manage Inventory** → **Computers**. Select the HMC name in the list and click the name. Select the Hardware tab and move the mouse over the HMC name. You see the device ID. In Figure 3-20 on page 87, the device ID for hmc1a is 11949.

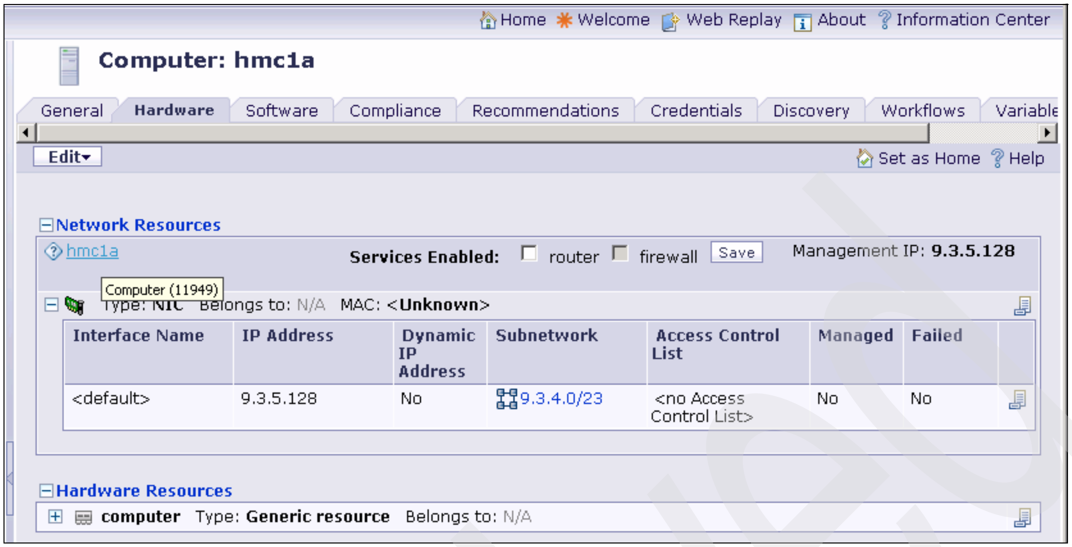

Computer: hmc1a

Network Resources

hmc1a   Services Enabled: □ router □ firewall   Management IP: **9.3.5.128**

Type: NIC   Belongs to: N/A   MAC: <Unknown>

Interface Name	IP Address	Dynamic IP Address	Subnetwork	Access Control List	Managed	Failed
<default>	9.3.5.128	No	9.3.4.0/23	<no Access Control List>	No	No

Hardware Resources

computer   Type: **Generic resource**   Belongs to: N/A

Figure 3-20   Determine the device ID of HMC in Hardware tab

After determining the device ID of HMC, send the command that was executed on HMC directly, as follows:

1. In the ExecuteCommand field, enter the `ls` command.
2. In the CredentialKey field, enter the value you specified, as shown in Figure 3-21.

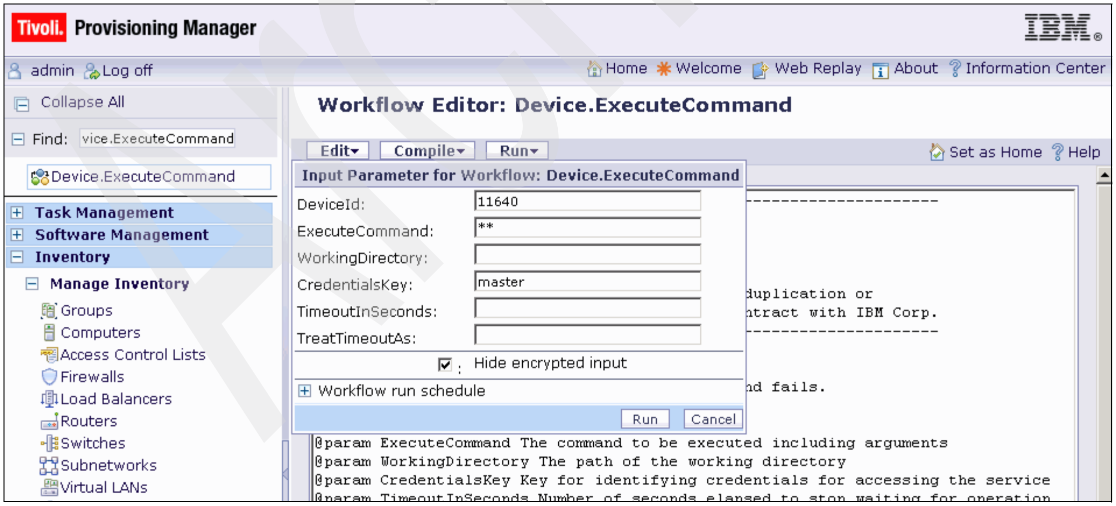

Workflow Editor: Device.ExecuteCommand

Input Parameter for Workflow: Device.ExecuteCommand

DeviceId:	11640
ExecuteCommand:	**
WorkingDirectory:	
CredentialsKey:	master
TimeoutInSeconds:	
TreatTimeoutAs:	

☑ Hide encrypted input

Workflow run schedule

```
---------------------
duplication or
ntract with IBM Corp.
---------------------

nd fails.
@param ExecuteCommand The command to be executed including arguments
@param WorkingDirectory The path of the working directory
@param CredentialsKey Key for identifying credentials for accessing the service
@param TimeoutInSeconds Number of seconds elapsed to stop waiting for operation
```

Figure 3-21   Run the ls command on HMC directly through TPM

After successfully sending the command to HMC from TPM, we see the results returned from HMC, as shown in Figure 3-22.

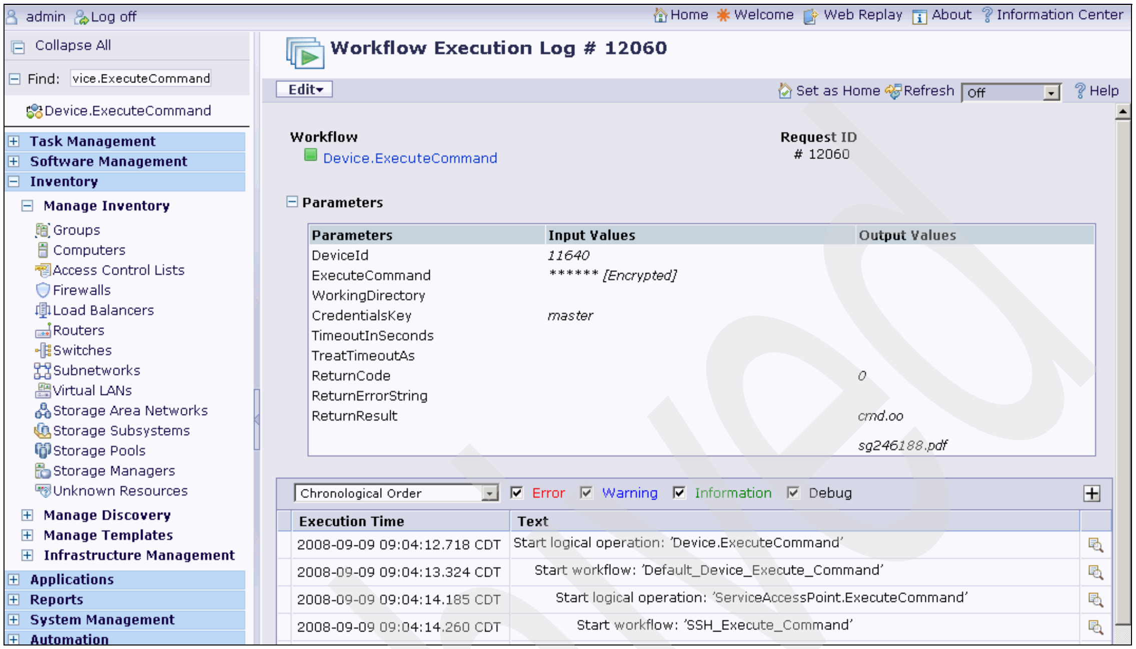

Workflow Execution Log # 12060

Workflow: Device.ExecuteCommand

Request ID: # 12060

Parameters

Parameters	Input Values	Output Values
DeviceId	*11640*	
ExecuteCommand	****** *[Encrypted]*	
WorkingDirectory		
CredentialsKey	*master*	
TimeoutInSeconds		
TreatTimeoutAs		
ReturnCode		*0*
ReturnErrorString		
ReturnResult		*cmd.oo* *sg246188.pdf*

Chronological Order ☑ Error ☑ Warning ☑ Information ☑ Debug

Execution Time	Text
2008-09-09 09:04:12.718 CDT	Start logical operation: 'Device.ExecuteCommand'
2008-09-09 09:04:13.324 CDT	Start workflow: 'Default_Device_Execute_Command'
2008-09-09 09:04:14.185 CDT	Start logical operation: 'ServiceAccessPoint.ExecuteCommand'
2008-09-09 09:04:14.260 CDT	Start workflow: 'SSH_Execute_Command'

Figure 3-22 Result of ls command returned from HMC to Workflow Execution Log in TPM

Creating and discovering a Virtual I/O Server

Create a Virtual I/O Server on the server you want to manage with TPM. Virtual I/O Server uses a restricted shell for security, so TPM cannot create a Virtual I/O Server yet. We must prepare and install the Virtual I/O Server partition manually for resource allocation before deployment. We also have to create one Virtual SCSI client partition profile and allocate virtual storage to that partition. After creating the Virtual I/O Server and the client partition, we can see the results, as shown in Figure 3-23 on page 89.

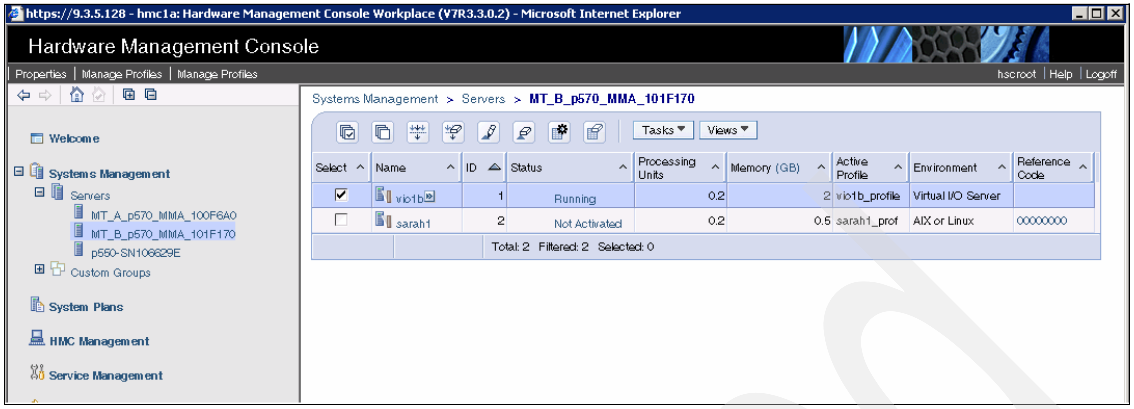

Select	Name	ID	Status	Processing Units	Memory (GB)	Active Profile	Environment	Reference Code
☑	vio1b	1	Running	0.2	2	vio1b_profile	Virtual I/O Server	
☐	sarah1	2	Not Activated	0.2	0.5	sarah1_prof	AIX or Linux	00000000

Total: 2 Filtered: 2 Selected: 0

Figure 3-23 Predefined Virtual I/O Server and logical partition on HMC

> **Note:** At the time of this writing, we can manage a Virtual I/O Server with TPM, but we cannot create a new Virtual I/O Server automatically with TPM.

After creating a Virtual I/O Server, discover that Virtual I/O Server by selecting **Inventory** → **Manage Discovery** → **Initial Discovery** → **New discovery**, as shown in Figure 3-24.

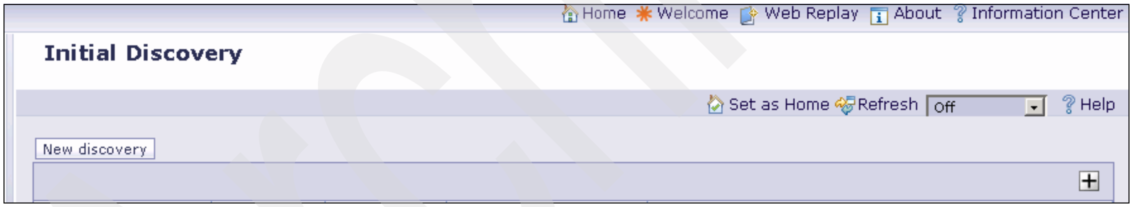

Figure 3-24 Initial Discovery

Enter the IP address of the Virtual I/O Server, and credential information, as shown in Figure 3-25 on page 90, and then click **Finish**.

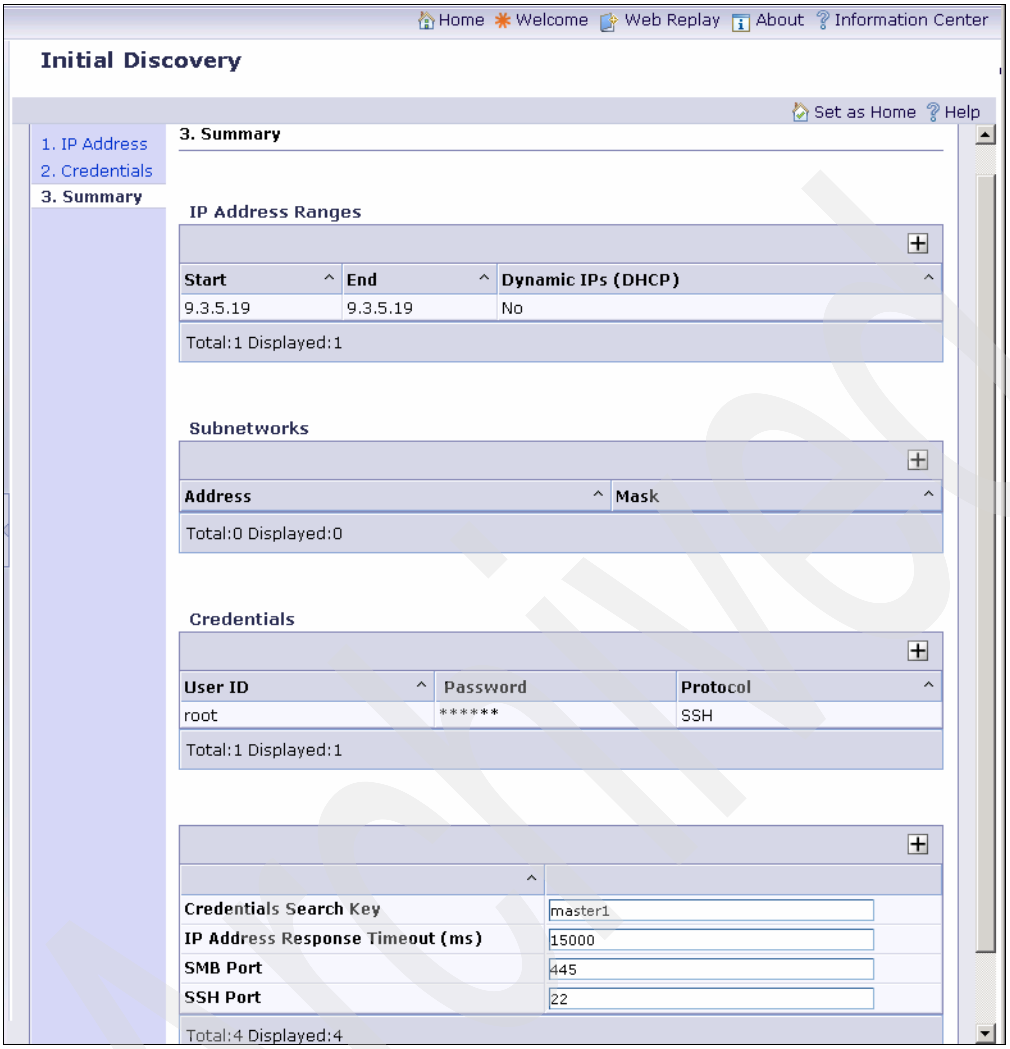

Initial Discovery

3. Summary

IP Address Ranges

Start	End	Dynamic IPs (DHCP)
9.3.5.19	9.3.5.19	No

Total:1 Displayed:1

Subnetworks

Address	Mask

Total:0 Displayed:0

Credentials

User ID	Password	Protocol
root	******	SSH

Total:1 Displayed:1

Credentials Search Key	master1
IP Address Response Timeout (ms)	15000
SMB Port	445
SSH Port	22

Total:4 Displayed:4

Figure 3-25 Parameter setting for Initial discovery of Virtual I/O Server

After TPM discovers the Virtual I/O Server:

1. Select **Inventory** → **Manage Inventory** → **Computers**.
2. Select Virtual I/O Server name, then select the General tab.
3. Click **Run** → **Run Inventory Scan**, as shown in Figure 3-26 on page 91.

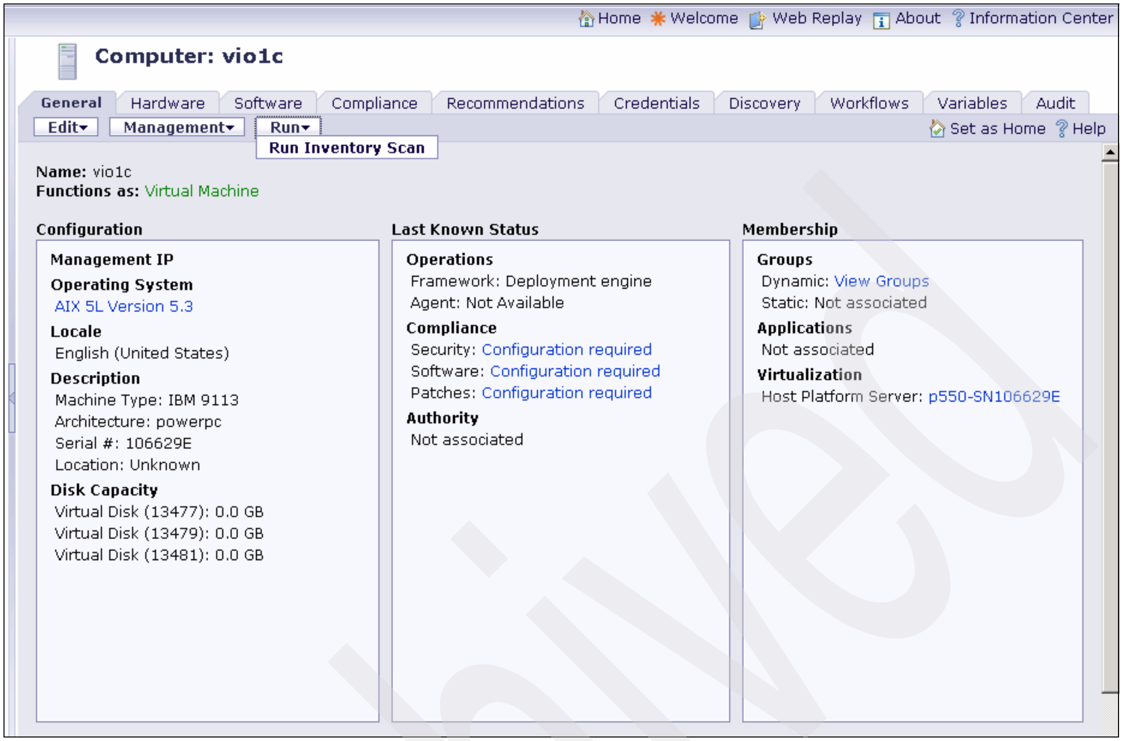

Computer: vio1c

Name: vio1c
Functions as: Virtual Machine

Configuration

Management IP

Operating System
AIX 5L Version 5.3

Locale
English (United States)

Description
Machine Type: IBM 9113
Architecture: powerpc
Serial #: 106629E
Location: Unknown

Disk Capacity
Virtual Disk (13477): 0.0 GB
Virtual Disk (13479): 0.0 GB
Virtual Disk (13481): 0.0 GB

Last Known Status

Operations
Framework: Deployment engine
Agent: Not Available

Compliance
Security: Configuration required
Software: Configuration required
Patches: Configuration required

Authority
Not associated

Membership

Groups
Dynamic: View Groups
Static: Not associated

Applications
Not associated

Virtualization
Host Platform Server: p550-SN106629E

Figure 3-26 Discover all Virtual I/O Server resource by Run Inventory Scan

All hardware and allocation information that Virtual I/O Server manages is displayed. See Figure 3-27 on page 92 and Figure 3-28 on page 93.

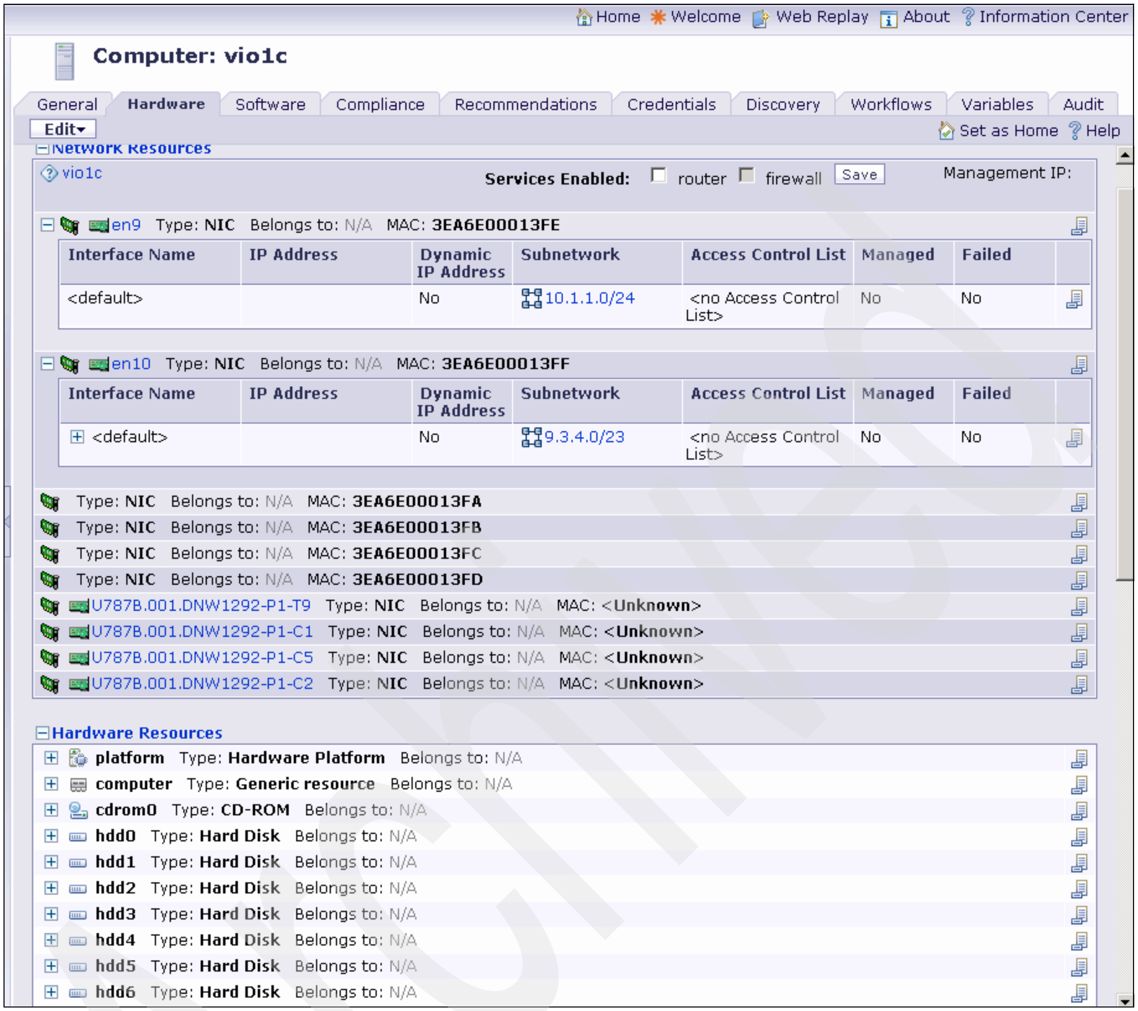

Computer: vio1c

Network Resources

vio1c **Services Enabled:** □ router □ firewall Save Management IP:

en9 Type: **NIC** Belongs to: N/A MAC: **3EA6E00013FE**

Interface Name	IP Address	Dynamic IP Address	Subnetwork	Access Control List	Managed	Failed
<default>		No	10.1.1.0/24	<no Access Control List>	No	No

en10 Type: **NIC** Belongs to: N/A MAC: **3EA6E00013FF**

Interface Name	IP Address	Dynamic IP Address	Subnetwork	Access Control List	Managed	Failed
⊞ <default>		No	9.3.4.0/23	<no Access Control List>	No	No

Type: **NIC** Belongs to: N/A MAC: **3EA6E00013FA**

Type: **NIC** Belongs to: N/A MAC: **3EA6E00013FB**

Type: **NIC** Belongs to: N/A MAC: **3EA6E00013FC**

Type: **NIC** Belongs to: N/A MAC: **3EA6E00013FD**

U787B.001.DNW1292-P1-T9 Type: **NIC** Belongs to: N/A MAC: <**Unknown**>

U787B.001.DNW1292-P1-C1 Type: **NIC** Belongs to: N/A MAC: <**Unknown**>

U787B.001.DNW1292-P1-C5 Type: **NIC** Belongs to: N/A MAC: <**Unknown**>

U787B.001.DNW1292-P1-C2 Type: **NIC** Belongs to: N/A MAC: <**Unknown**>

Hardware Resources

⊞ **platform** Type: **Hardware Platform** Belongs to: N/A

⊞ **computer** Type: **Generic resource** Belongs to: N/A

⊞ **cdrom0** Type: **CD-ROM** Belongs to: N/A

⊞ **hdd0** Type: **Hard Disk** Belongs to: N/A

⊞ **hdd1** Type: **Hard Disk** Belongs to: N/A

⊞ **hdd2** Type: **Hard Disk** Belongs to: N/A

⊞ **hdd3** Type: **Hard Disk** Belongs to: N/A

⊞ **hdd4** Type: **Hard Disk** Belongs to: N/A

⊞ **hdd5** Type: **Hard Disk** Belongs to: N/A

⊞ **hdd6** Type: **Hard Disk** Belongs to: N/A

Figure 3-27 Virtual I/O Server resource information

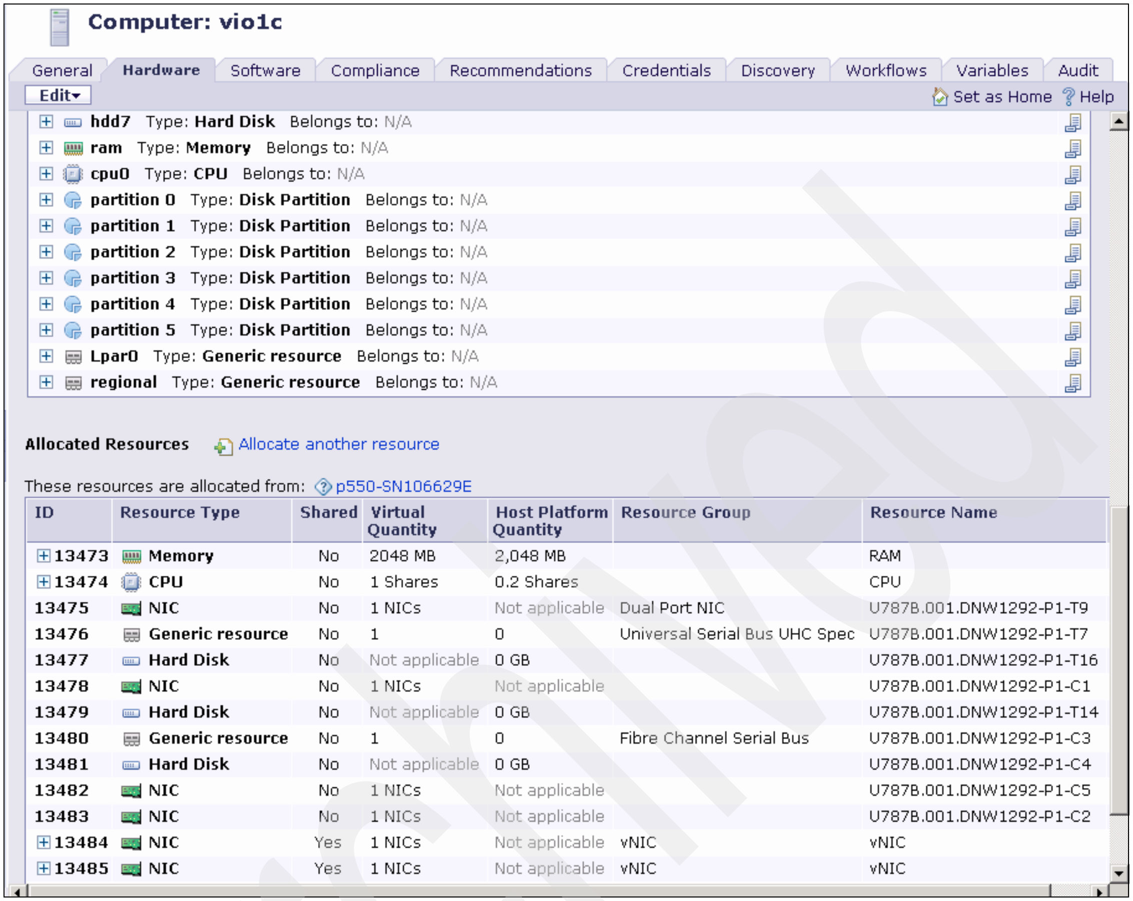

hdd7 Type: Hard Disk Belongs to: N/A
ram Type: Memory Belongs to: N/A
cpu0 Type: CPU Belongs to: N/A
partition 0 Type: Disk Partition Belongs to: N/A
partition 1 Type: Disk Partition Belongs to: N/A
partition 2 Type: Disk Partition Belongs to: N/A
partition 3 Type: Disk Partition Belongs to: N/A
partition 4 Type: Disk Partition Belongs to: N/A
partition 5 Type: Disk Partition Belongs to: N/A
Lpar0 Type: Generic resource Belongs to: N/A
regional Type: Generic resource Belongs to: N/A

Allocated Resources Allocate another resource

These resources are allocated from: p550-SN106629E

ID	Resource Type	Shared	Virtual Quantity	Host Platform Quantity	Resource Group	Resource Name
⊞13473	Memory	No	2048 MB	2,048 MB		RAM
⊞13474	CPU	No	1 Shares	0.2 Shares		CPU
13475	NIC	No	1 NICs	Not applicable	Dual Port NIC	U787B.001.DNW1292-P1-T9
13476	Generic resource	No	1	0	Universal Serial Bus UHC Spec	U787B.001.DNW1292-P1-T7
13477	Hard Disk	No	Not applicable	0 GB		U787B.001.DNW1292-P1-T16
13478	NIC	No	1 NICs	Not applicable		U787B.001.DNW1292-P1-C1
13479	Hard Disk	No	Not applicable	0 GB		U787B.001.DNW1292-P1-T14
13480	Generic resource	No	1	0	Fibre Channel Serial Bus	U787B.001.DNW1292-P1-C3
13481	Hard Disk	No	Not applicable	0 GB		U787B.001.DNW1292-P1-C4
13482	NIC	No	1 NICs	Not applicable		U787B.001.DNW1292-P1-C5
13483	NIC	No	1 NICs	Not applicable		U787B.001.DNW1292-P1-C2
⊞13484	NIC	Yes	1 NICs	Not applicable	vNIC	vNIC
⊞13485	NIC	Yes	1 NICs	Not applicable	vNIC	vNIC

Figure 3-28 Virtual I/O Server resource information (continued)

Discovering the managed system information

To discover system information, we must define HMC discovery configuration:

1. Select **Inventory** → **Manage Discovery** → **Discovery Configurations**.
2. Select the **Edit** tab and **Add Discovery Configuration**. The Configure Discovery panel is displayed. See Figure 3-29 on page 94.
3. For Discover, select **Other**; for Discovery Method, select **HMC Discovery**.

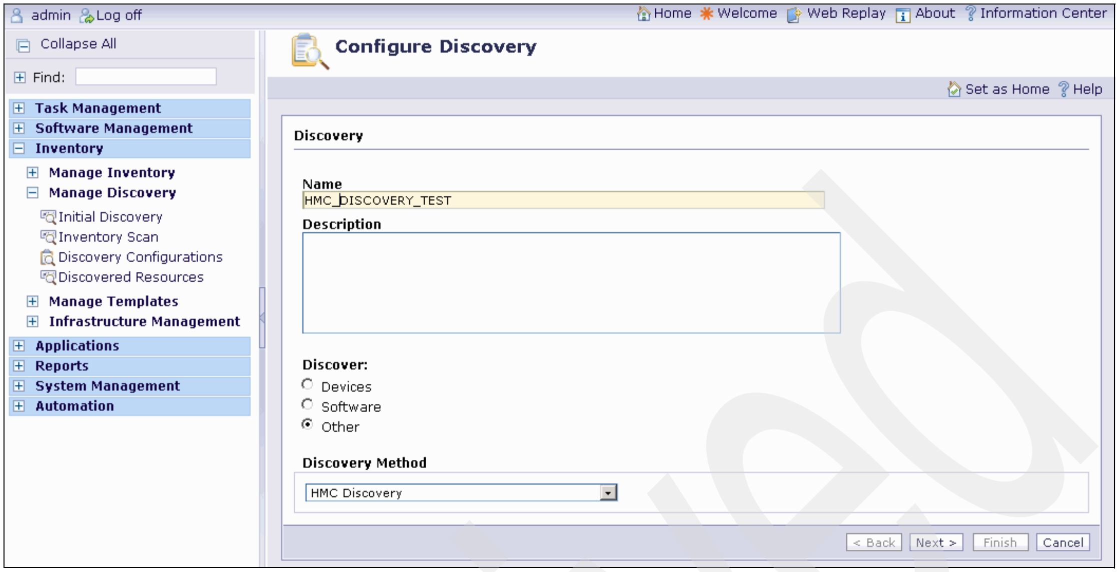

Name: HMC_DISCOVERY_TEST

Description:

Discover:
○ Devices
○ Software
● Other

Discovery Method: HMC Discovery

Figure 3-29 Configuring discovery method to discover system information

4. For the Computer (as shown in Figure 3-30 on page 95), specify the HMC that you previously defined.
5. For Central Electronics Complex (CEC), specify the CEC name you want to discover.

> **Important:** For CEC name, do not use any blanks. If you use blanks in the CEC name, TPM cannot discover the CEC.

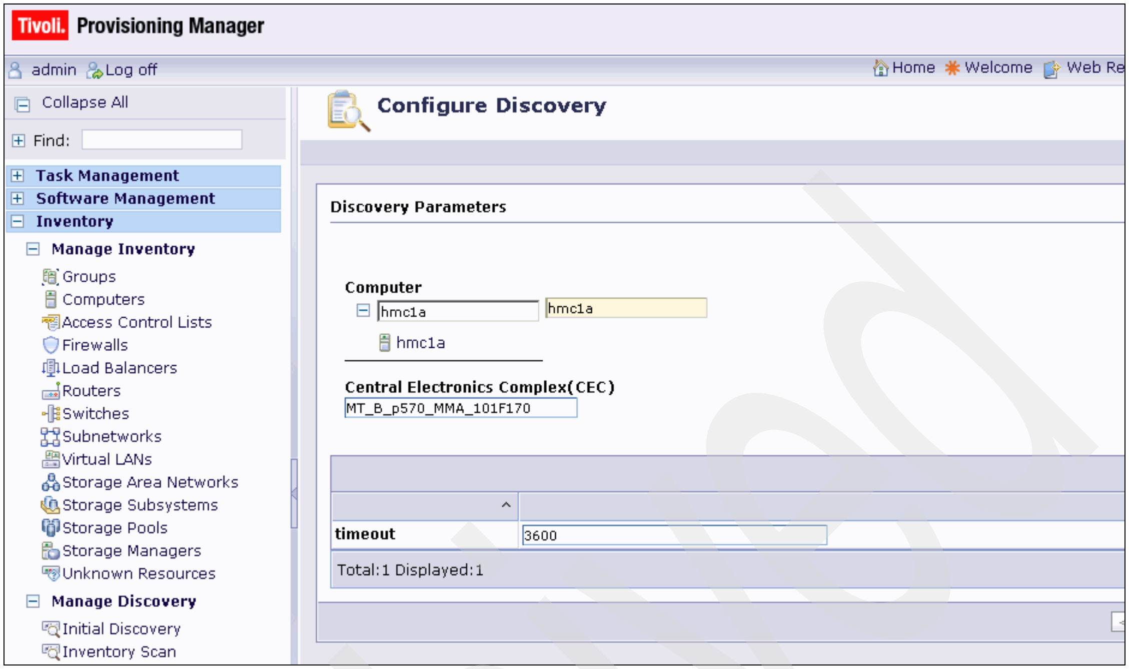

Configure Discovery

Discovery Parameters

Computer: hmc1a hmc1a
hmc1a

Central Electronics Complex(CEC): MT_B_p570_MMA_101F170

timeout: 3600

Total:1 Displayed:1

Figure 3-30 Configuring HMC and CEC name for discovery

6. After creating the discovery configuration, click **Run** (beside the configuration name).

   The Track Tasks panel displays, as shown in Figure 3-31 on page 96.

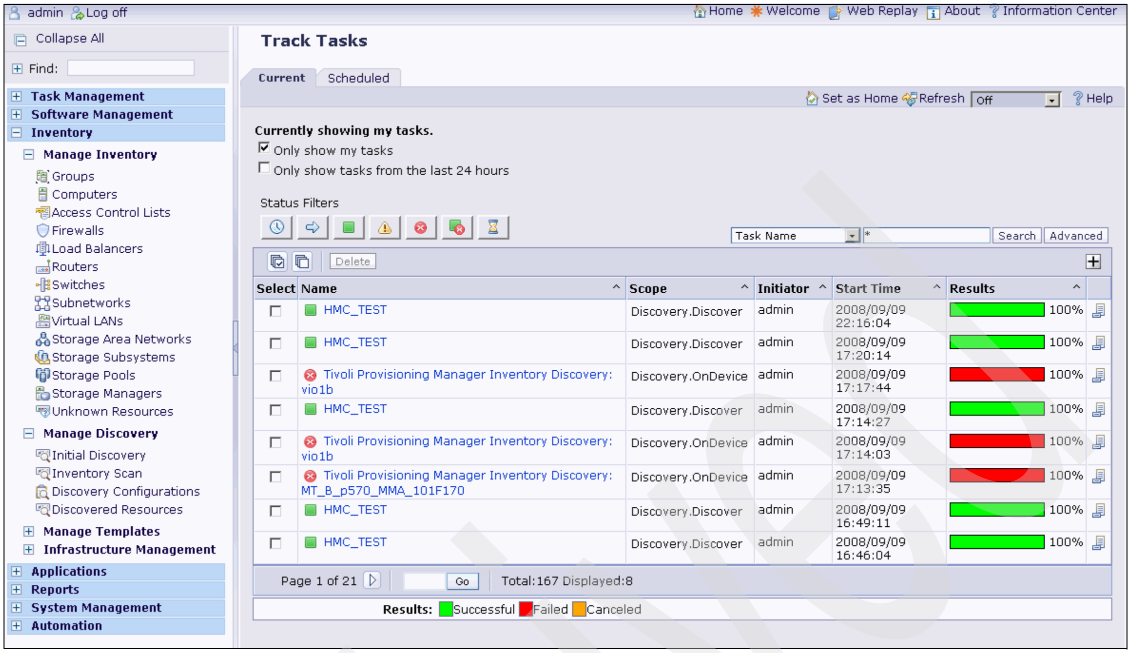

admin Log off Home Welcome Web Replay About Information Center

Collapse All

Find:

- Task Management
- Software Management
- Inventory
  - Manage Inventory
    - Groups
    - Computers
    - Access Control Lists
    - Firewalls
    - Load Balancers
    - Routers
    - Switches
    - Subnetworks
    - Virtual LANs
    - Storage Area Networks
    - Storage Subsystems
    - Storage Pools
    - Storage Managers
    - Unknown Resources
  - Manage Discovery
    - Initial Discovery
    - Inventory Scan
    - Discovery Configurations
    - Discovered Resources
  - Manage Templates
  - Infrastructure Management
- Applications
- Reports
- System Management
- Automation

Track Tasks

Current | Scheduled

Set as Home Refresh Off Help

Currently showing my tasks.

☑ Only show my tasks

☐ Only show tasks from the last 24 hours

Status Filters

Task Name * Search Advanced

Delete

Select	Name	Scope	Initiator	Start Time	Results
☐	🟩 HMC_TEST	Discovery.Discover	admin	2008/09/09 22:16:04	🟩 100%
☐	🟩 HMC_TEST	Discovery.Discover	admin	2008/09/09 17:20:14	🟩 100%
☐	⊗ Tivoli Provisioning Manager Inventory Discovery: vio1b	Discovery.OnDevice	admin	2008/09/09 17:17:44	🟥 100%
☐	🟩 HMC_TEST	Discovery.Discover	admin	2008/09/09 17:14:27	🟩 100%
☐	⊗ Tivoli Provisioning Manager Inventory Discovery: vio1b	Discovery.OnDevice	admin	2008/09/09 17:14:03	🟥 100%
☐	⊗ Tivoli Provisioning Manager Inventory Discovery: MT_B_p570_MMA_101F170	Discovery.OnDevice	admin	2008/09/09 17:13:35	🟥 100%
☐	🟩 HMC_TEST	Discovery.Discover	admin	2008/09/09 16:49:11	🟩 100%
☐	🟩 HMC_TEST	Discovery.Discover	admin	2008/09/09 16:46:04	🟩 100%

Page 1 of 21 Go Total:167 Displayed:8

Results: 🟩Successful 🟥Failed 🟧Canceled

Figure 3-31 Track Tasks displays the process of discovery task

7. After HMC discovery, view the CEC information and LPAR information by selecting **Inventory** → **Manage Inventory** → **Computers**. In the Hardware tab, view all the Network Resources, Hardware Resources, and Resource Allocations, as shown in Figure 3-32 on page 97.

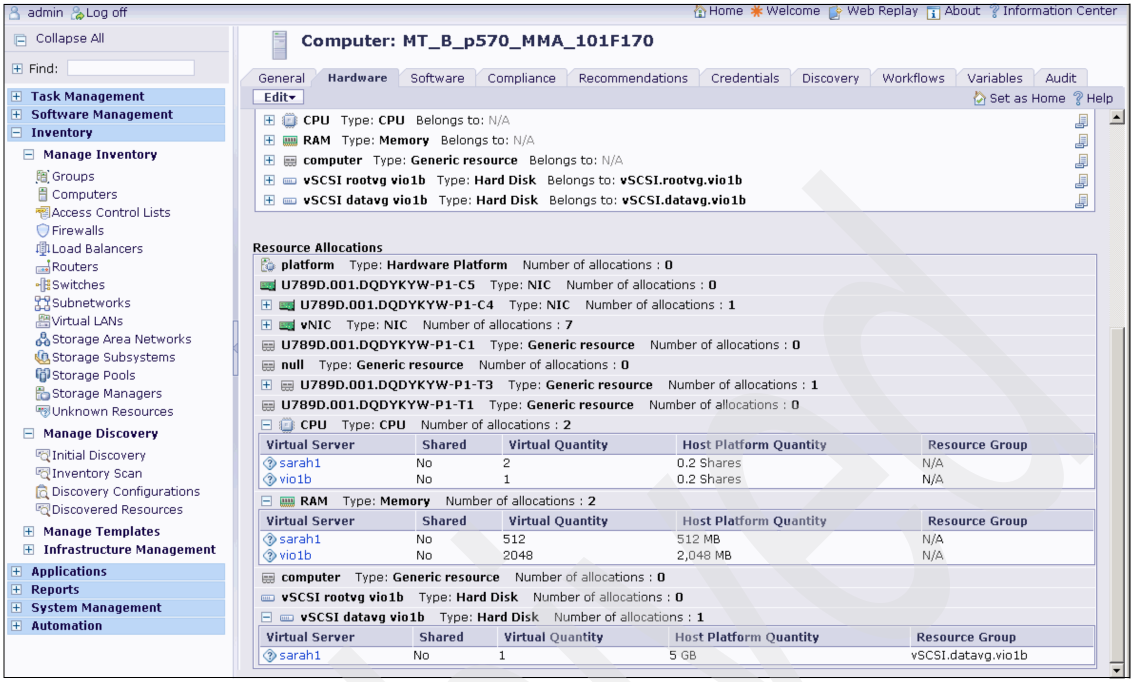

Computer: MT_B_p570_MMA_101F170

- CPU Type: CPU Belongs to: N/A
- RAM Type: Memory Belongs to: N/A
- computer Type: Generic resource Belongs to: N/A
- vSCSI rootvg vio1b Type: Hard Disk Belongs to: vSCSI.rootvg.vio1b
- vSCSI datavg vio1b Type: Hard Disk Belongs to: vSCSI.datavg.vio1b

Resource Allocations

- platform Type: Hardware Platform Number of allocations : 0
- U789D.001.DQDYKYW-P1-C5 Type: NIC Number of allocations : 0
- U789D.001.DQDYKYW-P1-C4 Type: NIC Number of allocations : 1
- vNIC Type: NIC Number of allocations : 7
- U789D.001.DQDYKYW-P1-C1 Type: Generic resource Number of allocations : 0
- null Type: Generic resource Number of allocations : 0
- U789D.001.DQDYKYW-P1-T3 Type: Generic resource Number of allocations : 1
- U789D.001.DQDYKYW-P1-T1 Type: Generic resource Number of allocations : 0
- CPU Type: CPU Number of allocations : 2

Virtual Server	Shared	Virtual Quantity	Host Platform Quantity	Resource Group
sarah1	No	2	0.2 Shares	N/A
vio1b	No	1	0.2 Shares	N/A

- RAM Type: Memory Number of allocations : 2

Virtual Server	Shared	Virtual Quantity	Host Platform Quantity	Resource Group
sarah1	No	512	512 MB	N/A
vio1b	No	2048	2,048 MB	N/A

- computer Type: Generic resource Number of allocations : 0
- vSCSI rootvg vio1b Type: Hard Disk Number of allocations : 0
- vSCSI datavg vio1b Type: Hard Disk Number of allocations : 1

Virtual Server	Shared	Virtual Quantity	Host Platform Quantity	Resource Group
sarah1	No	1	5 GB	vSCSI.datavg.vio1b

Figure 3-32 All logical partitions and resource information discovered under CEC

8. Make sure that all the resources have **Managed** and **Partitionable** boxes checked in the Hardware Resources panel. Beside each resource, right-click the icon and select **Properties** as shown in Figure 3-33.

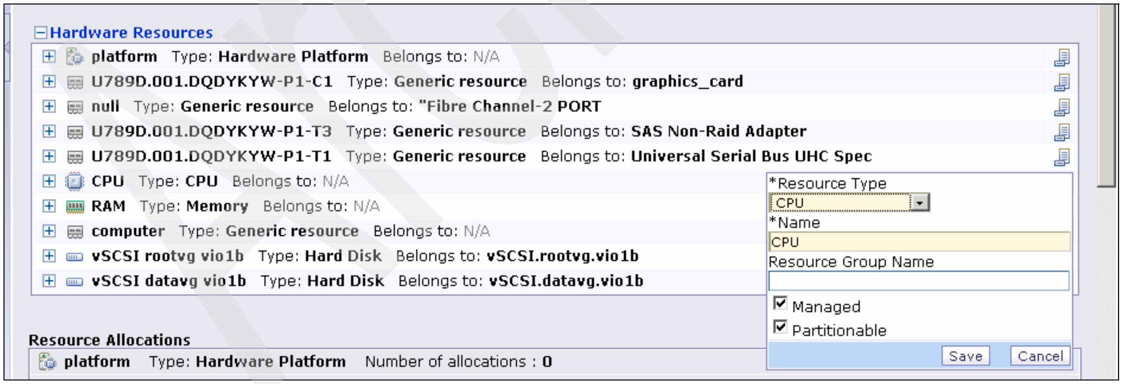

Hardware Resources

- platform Type: Hardware Platform Belongs to: N/A
- U789D.001.DQDYKYW-P1-C1 Type: Generic resource Belongs to: graphics_card
- null Type: Generic resource Belongs to: "Fibre Channel-2 PORT
- U789D.001.DQDYKYW-P1-T3 Type: Generic resource Belongs to: SAS Non-Raid Adapter
- U789D.001.DQDYKYW-P1-T1 Type: Generic resource Belongs to: Universal Serial Bus UHC Spec
- CPU Type: CPU Belongs to: N/A
- RAM Type: Memory Belongs to: N/A
- computer Type: Generic resource Belongs to: N/A
- vSCSI rootvg vio1b Type: Hard Disk Belongs to: vSCSI.rootvg.vio1b
- vSCSI datavg vio1b Type: Hard Disk Belongs to: vSCSI.datavg.vio1b

*Resource Type: CPU

*Name: CPU

Resource Group Name:

☑ Managed

☑ Partitionable

Resource Allocations

- platform Type: Hardware Platform Number of allocations : 0

Figure 3-33 Check resource status as managed and partitionable

Creating a virtual server template

Create a virtual server template to specify the amount and types of resources that we want for the allocation of a virtual server.

Select **Inventory** → **Manage Templates** → **Virtual Storage Templates**. Copy `Power5_micro_partition` as a new file name, and edit that file appropriately.

In the example, as shown in Figure 3-34, CPU capacity is 0.2 Shares (Host Platform Quantity) of physical capacity and virtual processor number is 2 Shares (Virtual Quantity). The memory is 512 MB and hard disk is allocated 20.0 GB in `datavg` of Virtual I/O Server `vio1b`. We allocated 2 vNICs, one for OS installation and management and one for users.

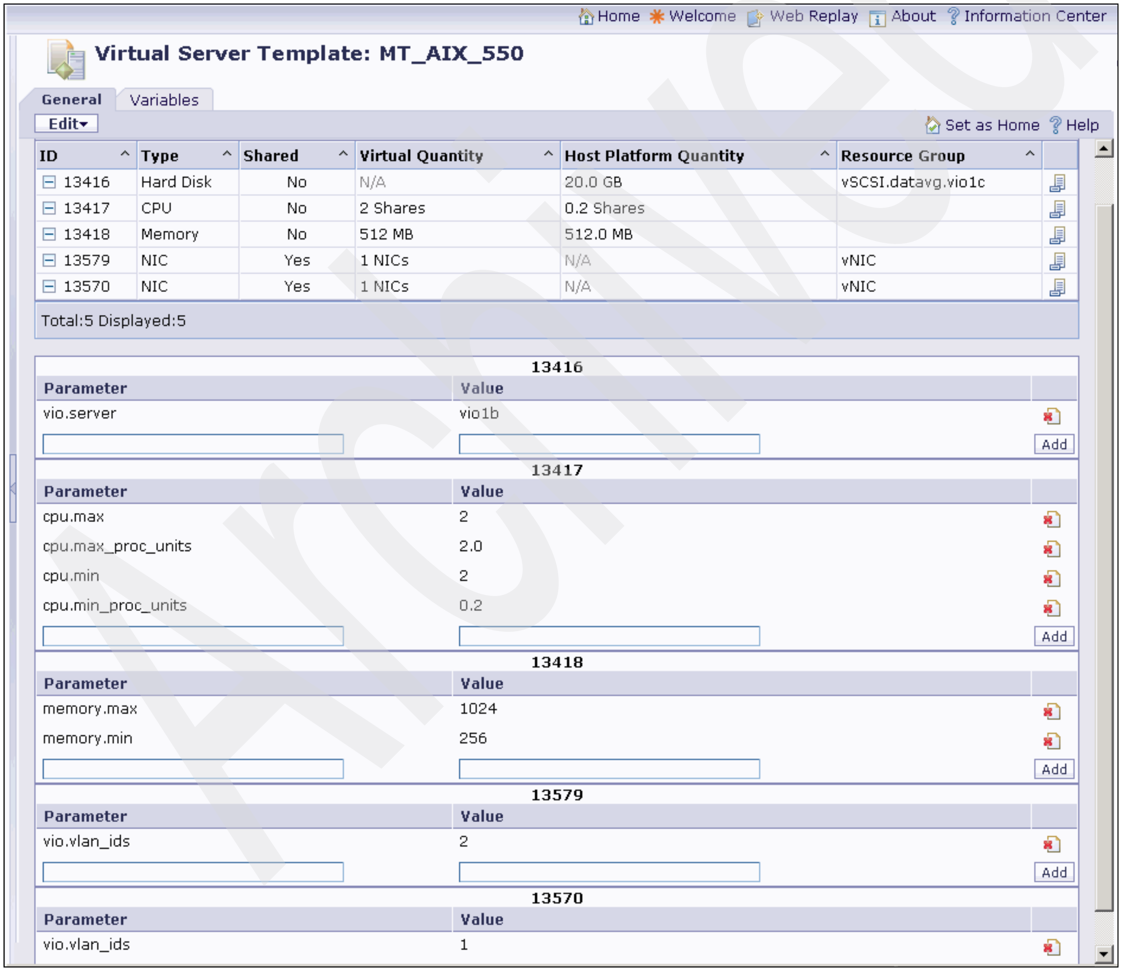

Virtual Server Template: MT_AIX_550

ID	Type	Shared	Virtual Quantity	Host Platform Quantity	Resource Group
⊟ 13416	Hard Disk	No	N/A	20.0 GB	vSCSI.datavg.vio1c
⊟ 13417	CPU	No	2 Shares	0.2 Shares	
⊟ 13418	Memory	No	512 MB	512.0 MB	
⊟ 13579	NIC	Yes	1 NICs	N/A	vNIC
⊟ 13570	NIC	Yes	1 NICs	N/A	vNIC

Total:5 Displayed:5

13416

Parameter	Value
vio.server	vio1b

13417

Parameter	Value
cpu.max	2
cpu.max_proc_units	2.0
cpu.min	2
cpu.min_proc_units	0.2

13418

Parameter	Value
memory.max	1024
memory.min	256

13579

Parameter	Value
vio.vlan_ids	2

13570

Parameter	Value
vio.vlan_ids	1

Figure 3-34 Virtual server template customized for environment

Now, TPM can manage LPARs by HMC and the host platform. On the host platform, we can create an LPAR and manage it with HMC and TPM. To create the LPAR, TPM must allocate the IP address and host name in DNS server, as shown in Figure 3-35.

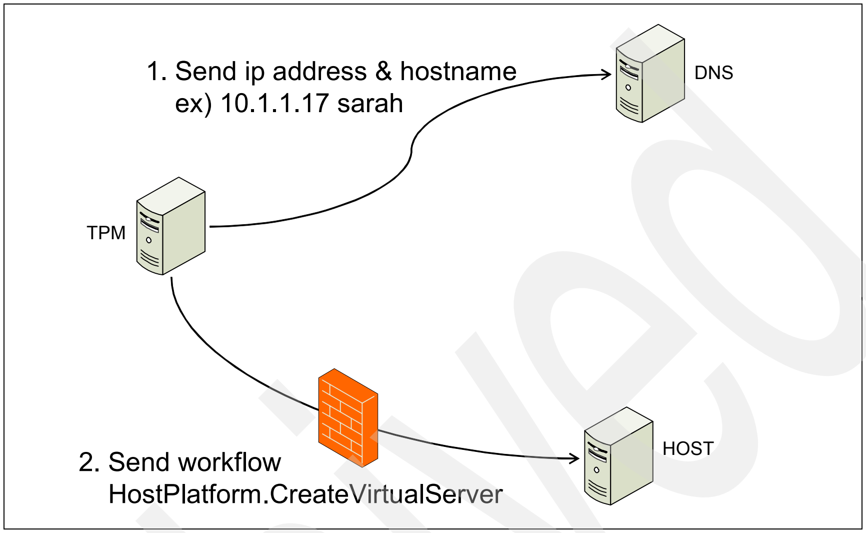
1. Send ip address & hostname
ex) 10.1.1.17 sarah
DNS
TPM
HOST
2. Send workflow
HostPlatform.CreateVirtualServer

Figure 3-35 Creating logical partition after DNS registration

To register the IP address and host name, we used a shell script, as shown in Example 3-2 on page 100.

Example 3-2 Shell script that register host name and IP address in DNS

```
#!/usr/bin/ksh

DNS_SERVER=itm1a.austin.ibm.com
TMP_FILE=/tmp/tmp.data

if [ $1 ]
then
        rsh $DNS_SERVER -l root cat /var/named/named.data > $TMP_FILE
        echo "$1         9999999 IN      A       10.1.1.$2" >> $TMP_FILE
        rcp $TMP_FILE $DNS_SERVER:/var/named/named.data
        rm $TMP_FILE

        rsh $DNS_SERVER -l root cat /var/named/named.rev > $TMP_FILE
        echo "$2     IN PTR $1.austin.ibm.com." >> $TMP_FILE
        rcp $TMP_FILE $DNS_SERVER:/var/named/named.rev
        rm $TMP_FILE

        rsh $DNS_SERVER -l root refresh -s named

else
        echo "Usage:dns_register.sh [hostname] [ip address]"
        echo "Example) dns_register sarah 17"
fi
```

Now we can determine host name in DNS server. The result are shown in Example 3-3.

Example 3-3 The result of Example 3-2

```
tpm1a:/tmp#>./dns_register.sh sarah 17
0513-095 The request for subsystem refresh was completed successfully.
tpm1a:/tmp#>nslookup
Default Server:  itm1a.austin.ibm.com
Address:  9.3.5.10

> sarah
Server:  itm1a.austin.ibm.com
Address:  9.3.5.10

Name:    sarah.austin.ibm.com
Address:  10.1.1.17

> exit
```

Allocating a new virtual server

To allocate the virtual server or LPAR in the TPM Web interface:

1. Select the name of the managed computer or host platform computer.
2. Select **Edit** → **Allocate a Virtual Server**.
3. Enter the partition's name and the select template, as shown in Figure 3-36.
4. Click **Save**.

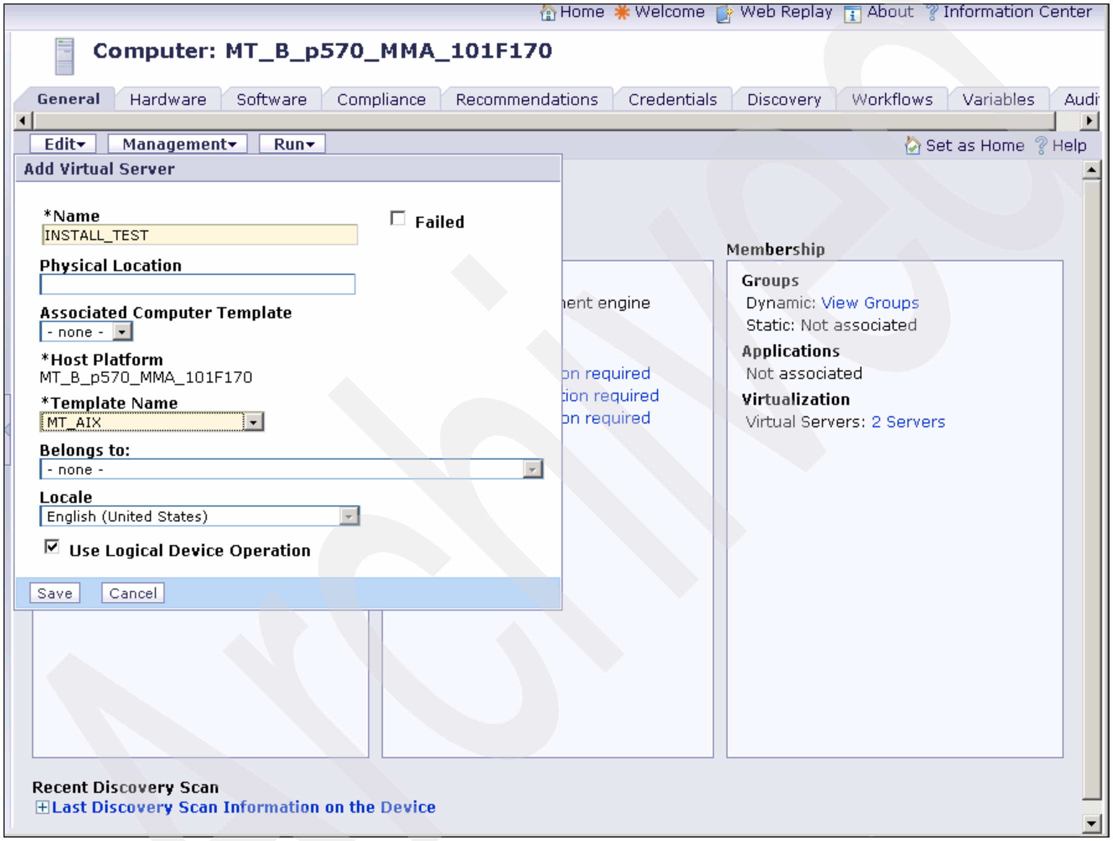

Computer: MT_B_p570_MMA_101F170

Add Virtual Server

*Name: INSTALL_TEST

□ Failed

Physical Location:

Associated Computer Template: - none -

*Host Platform: MT_B_p570_MMA_101F170

*Template Name: MT_AIX

Belongs to: - none -

Locale: English (United States)

☑ Use Logical Device Operation

ent engine

on required
ion required
on required

Membership

Groups
Dynamic: View Groups
Static: Not associated

Applications
Not associated

Virtualization
Virtual Servers: 2 Servers

Recent Discovery Scan
⊞ Last Discovery Scan Information on the Device

Figure 3-36 Add virtual server with template under CEC

The virtual server create workflow starts as shown in Figure 3-37 on page 102.

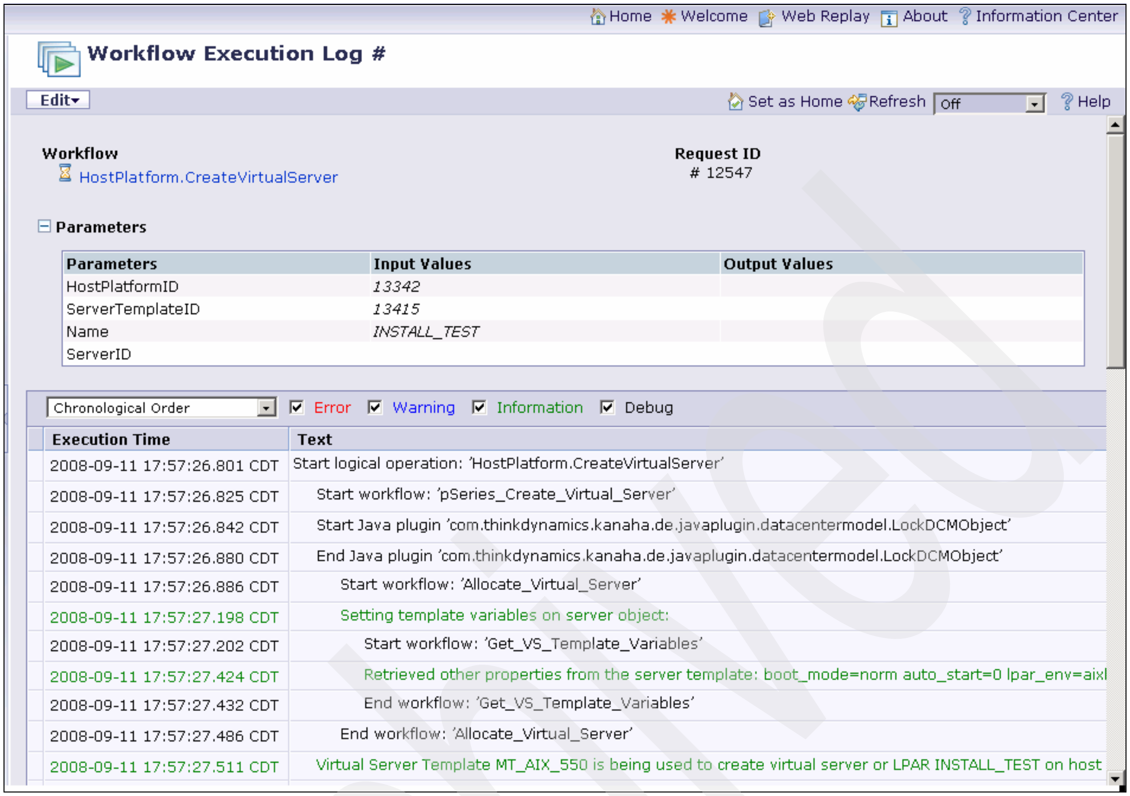

Workflow Execution Log #

Workflow
HostPlatform.CreateVirtualServer

Request ID
12547

⊟ Parameters

Parameters	Input Values	Output Values
HostPlatformID	*13342*	
ServerTemplateID	*13415*	
Name	*INSTALL_TEST*	
ServerID		

Chronological Order ☑ Error ☑ Warning ☑ Information ☑ Debug

Execution Time	Text
2008-09-11 17:57:26.801 CDT	Start logical operation: 'HostPlatform.CreateVirtualServer'
2008-09-11 17:57:26.825 CDT	Start workflow: 'pSeries_Create_Virtual_Server'
2008-09-11 17:57:26.842 CDT	Start Java plugin 'com.thinkdynamics.kanaha.de.javaplugin.datacentermodel.LockDCMObject'
2008-09-11 17:57:26.880 CDT	End Java plugin 'com.thinkdynamics.kanaha.de.javaplugin.datacentermodel.LockDCMObject'
2008-09-11 17:57:26.886 CDT	Start workflow: 'Allocate_Virtual_Server'
2008-09-11 17:57:27.198 CDT	Setting template variables on server object:
2008-09-11 17:57:27.202 CDT	Start workflow: 'Get_VS_Template_Variables'
2008-09-11 17:57:27.424 CDT	Retrieved other properties from the server template: boot_mode=norm auto_start=0 lpar_env=aixl
2008-09-11 17:57:27.432 CDT	End workflow: 'Get_VS_Template_Variables'
2008-09-11 17:57:27.486 CDT	End workflow: 'Allocate_Virtual_Server'
2008-09-11 17:57:27.511 CDT	Virtual Server Template MT_AIX_550 is being used to create virtual server or LPAR INSTALL_TEST on host

Figure 3-37 The Workflow Execution Log of creating virtual server

For more information:

- See *Deployment Guide Series: IBM Tivoli Provisioning Manager Version 5.1*, SG24-7261.
- See the product documentation:

  http://publib.boulder.ibm.com/infocenter/tivihelp/v20r1/index.jsp

If you find that clicking the mouse numerous times is inconvenient, you may use the *Web Replay* function that is provided by TPM. For more information, see "Web Replay" on page 128.

3.5.4 Allocation with workload partition

The Workload Partitions (WPARs) provide a solution for highly virtualized operating system environments. This virtualization feature was introduced in AIX V6.1, and no separate software license is required.

In a WPAR environment, each WPAR instance can run a different workload completely isolated from the other WPARs that are running in the same LPAR. Each WPAR can have its own applications. See 2.1.3, "AIX 6 workload partitions" on page 25 for more information.

In a multitenant environment, a single tenant can also be a company that is providing services for multiple subtenants. This tenant can use the WPAR to create isolated environments for each subtenants, in the same AIX global environment. Figure 3-38 shows the various configurations with WPAR.

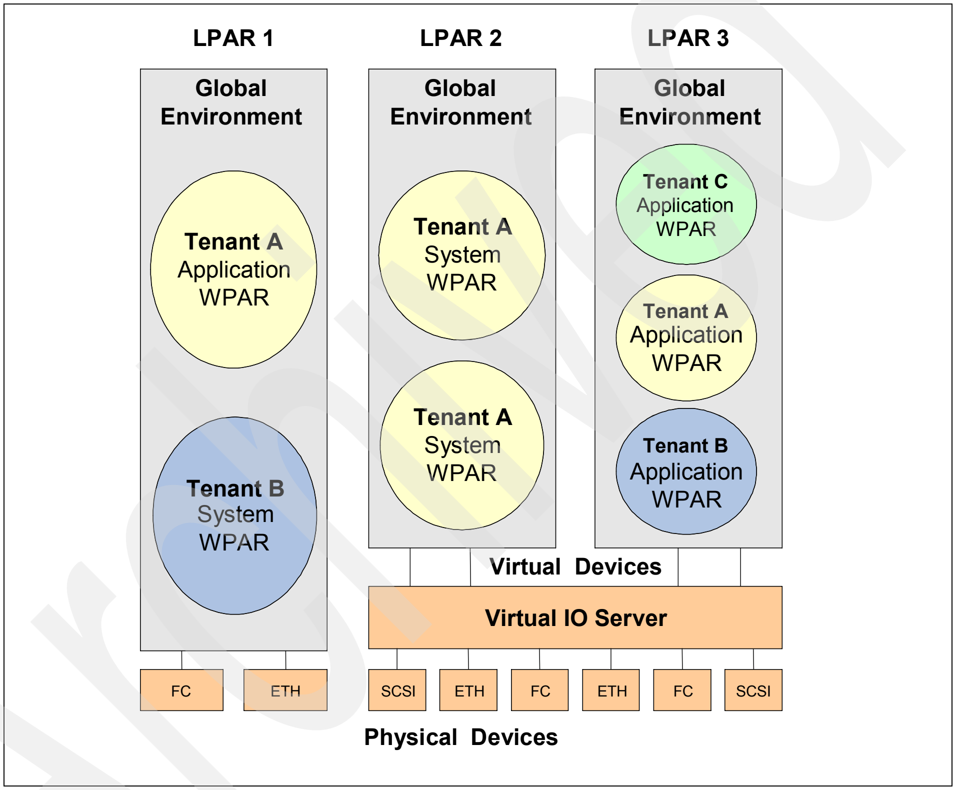

Figure 3-38 Examples of workload partitions

To create any type of WPAR, use either the **`smit wpar`** command to launch the dialog for WPAR, which can save time, or use the command line to enter one command at a time to do the configuration. Commands are available to create a system WPAR or an application WPAR.

In the environment used in this book, we used the AIX V6.1 command line to create and to configure a system WPAR called *TenantA* and an application WPAR called *TenantB*.

Managing system WPAR

This section is an overview of how to create, configure, and start a simple system WPAR, as we used in our lab environment.

Creating the system WPAR

Creating a simple system WPAR by using the **mkwpar** command is shown in Example 3-4.

Example 3-4 Creating a system WPAR by using the command line interface

```
$ mkwpar -n TenantA

mkwpar: Creating file systems...
 /
 /home
 /opt
 /proc
 /tmp
 /usr
 /var
Mounting all workload partition file systems.
x ./usr
x ./lib
x ./admin
x ./admin/tmp
x ./audit
x ./dev
x ./etc
..
.syncroot: Processing root part installation status.
syncroot: Synchronizing installp software.
+---------------------------------------------------------------------+
                    Pre-installation Verification...
+---------------------------------------------------------------------+
Verifying selections...done
Verifying requisites...done
Results...

SUCCESSES
---------
  Filesets listed in this section passed pre-installation verification
  and will be installed.

  Selected Filesets
  -----------------
  Java5.sdk 5.0.0.160                         # Java SDK 32-bit
...
Installation Summary
```

```
--------------------
Name                        Level           Part        Event       Result
-----------------------------------------------------------------------
bos.rte.install             6.1.1.1         ROOT        COMMIT      SUCCESS
syncroot: Processing root part installation status.
syncroot: Installp root packages are currently synchronized.
syncroot: RPM root packages are currently synchronized.
syncroot: Root part is currently synchronized.
syncroot: Returns Status = SUCCESS
Workload partition TenantB created successfully.
mkwpar: 0960-390 To start the workload partition, execute the following as root: startwpar [-v]
TenantA
```

Listing the created system WPAR by using the **lswpar** command is shown in Example 3-5.

Example 3-5 Listing the system WPARs

```
$ lswpar
Name       State  Type  Hostname   Directory
----------------------------------------------------
TenantA  D      S     TenantA  /wpars/TenantA
```

So far, the WPAR was created without any IP information and the status (`State`) is down (`D`). We need to set an IP address that is in the same subnet as the Global Environment interface. Now we are configuring the IP address to be assigned for this system WPAR.

Configuring the system WPAR

Setting network parameters is shown in Example 3-6. The example shows the **ifconfig** command output from the global environment and the **chwpar** command to create the network interface address in the WPAR.

Example 3-6 Setting network parameters

```
$ ifconfig en6
en6:
flags=1e080863,180<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,GR
OUPRT,64BIT,CHECKSUM_OFFLOAD(ACTIVE),CHAIN>
        inet 10.1.60.1 netmask 0xffffff00 broadcast 10.1.60.255
         tcp_sendspace 262144 tcp_recvspace 262144 rfc1323 1
$ chwpar -N interface=en6 address=10.1.60.2 netmask=255.255.255.0
TenantA
```

Starting the system WPAR

After configuring the IP address, we can start the WPAR. Starting and verifying the WPAR is shown in Example 3-7.

Example 3-7 Show the startwpar and verification tasks for the TenantA WPAR

```
$ startwpar TenantA
Starting workload partition 'TenantA'.
Mounting all workload partition file systems.
Loading workload partition.
Exporting workload partition devices.
Starting workload partition subsystem 'cor_TenantA'.
0513-059 The cor_TenantA Subsystem has been started. Subsystem PID is 712924.
Verifying workload partition startup.

$ ifconfig en6
en6:
flags=1e080863,180<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64
BIT,CHECKSUM_OFFLOAD(ACTIVE),CHAIN>
        inet 10.1.60.1 netmask 0xffffff00 broadcast 10.1.60.255
        inet 10.1.60.2 netmask 0xffffff00 broadcast 10.1.60.255
         tcp_sendspace 262144 tcp_recvspace 262144 rfc1323 1

$ ping 10.1.60.2
PING 10.1.60.2 (10.1.60.2): 56 data bytes
64 bytes from 10.1.60.2: icmp_seq=0 ttl=255 time=0 ms
64 bytes from 10.1.60.2: icmp_seq=1 ttl=255 time=0 ms
...

$ clogin TenantA
***********************************************************************
*
*  Welcome to AIX Version 6.1!

*                                                                             *
*  Please see the README file in /usr/lpp/bos for information pertinent * to
this release of the AIX Operating System.                                  *

*                                                                             *
***********************************************************************

# hostname
TenantA
```

The system WPAR is now up and running and the status (`State`) was changed to active (`A`) as shown in Example 3-8 on page 107.

Example 3-8 lswpar command showing the WPAR active

```
$ lswpar
Name       State  Type  Hostname   Directory
----------------------------------------------------
TenantA  A      S     TenantA  /wpars/TenantA
```

Managing application WPAR

The application WPAR only exists while the script or application that is associated with the WPAR is running. After the script or application is done, the WPAR is terminated. Before creating the application WPAR, we ensure that the application or script exists and has the correct permissions to be executed. For this test, a shell script, `/tmp/wpar/TenantB` file, as shown in Example 3-9, was created to provide the application WPAR executable.

Example 3-9 Script file used for the application WPAR

```
$ cat /tmp/wpar/TenantB
# !/bin/ksh
# Script TenantB
# This script was created to run the TenantB Application WPAR
#
# Created in 2008
# Objective : Test for WPAR environment
#
# Redirect the STDOUT and STDERR to a logfile
exec 1>/tmp/${0##*/}.log 2>&1
exec 2>&1
set -x

echo " #### Creating Empty files  #### "

for i in 1 2 3 4 5 6 7 8 9
do
touch /tmp/wpar/TenantB_file$i
done
```

Creating an application WPAR

To create the application WPAR use the **`wparexec`** command. The argument is the script or the application binary file as shown in Example 3-10 on page 108.

Example 3-10 The command wparexec runs the script under the TenantB WPAR

```
$ wparexec /tmp/wpar/TenantB

Starting workload partition 'TenantB'.
Mounting all workload partition file systems.
Loading workload partition.
Shutting down all workload partition processes.

$ ls
TenantB TenantB_file2  TenantB_file4  TenantB_file6 TenantB_file8
TenantB_file1  TenantB_file3  TenantB_file5  TenantB_file7
TenantB_file9
```

For this example we can say the application WPAR is not created but executed. And we can run several scripts or applications at the same time. All commands are executed from the global environment partition.

Fore more information about workload partitions, see the following sources:

- *Workload Partition Management in IBM AIX Version 6.1,* SG24-7656
- *Introduction to workload Partition Management in IBM AIX Version 6.1,* SG24-7431

3.5.5 Allocation with IBM Director

IBM Director product offers the possibility to manage, monitor, and allocate resources for complex and mixed environments.

IBM Director product can use the Virtual I/O Server to create virtualized LPARs or create LPARs by using dedicated resources.

To help you become more familiar with the product, we list several terms that are used in an IBM Director environment, to differentiate the equipment:

Management servers One or more servers where the IBM Director Server is installed

Managed system Any type of server or LPAR that is managed by the IBM Director

SNMP devices Network devices or systems that use the Simple Network Management Protocol (SNMP), to provide configuration data.

Managed objects Additional objects like a platform (POWER6 system) and chassis (for blade servers)

IBM Director major components

The IBM Director suite has various components to provide full administration control over the managed environment.

The Table 3-2 lists major components and a brief description of each component.

Table 3-2 Major IBM Director components

Component	Description
IBM Director Core Services	This is installed on a managed system to provide hardware-specific functions; used for management
IBM Director Agent	This is installed on a managed system to provide multiple functions as hardware and software operations.
IBM Director Console	This is installed on any machine that is able to provide a graphical user interface. The console is used to perform tasks in the IBM Director environment
IBM Director Server	This must be installed on the management server. It provides all the management information and functions of IBM Director environment. It tuns a database to keep the information about the managed systems. It is the main component.

The IBM Director management server connects to the managed systems or objects, and gathers information about the environment. The basic configuration is shown in Figure 3-39 on page 110. After that we are able to perform allocation tasks with the IBM Director console.

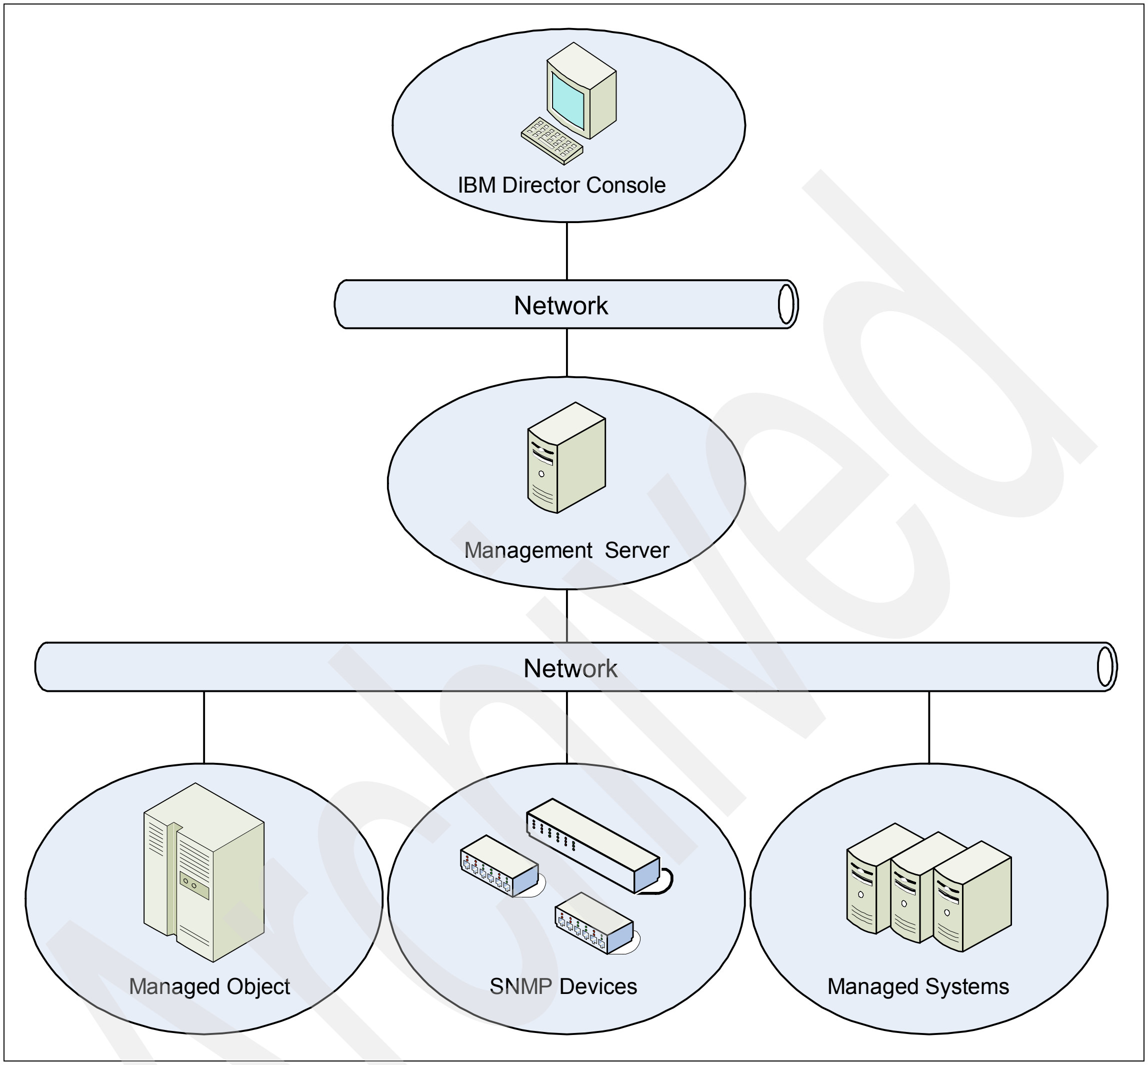

Figure 3-39   Basic IBM Director configuration and its components

The IBM Director Server must also have product extensions so that it can fully manage the Power Systems servers. Several available extensions are listed in Table 3-3 on page 111.

Table 3-3 IBM Director extensions that are used to manage a virtual environment

Extension name	Extension functionality
IBM Systems Director Virtualization Manager Server	The virtualization manager can be used to perform operations in a Virtual I/O Server environment. By using this feature, we can create a totally virtualized partition, creating virtual devices in the Virtual I/O Server and mapping them to the LPAR. All tasks can be performed in the IBM Director console.
Virtualization Manager Agent for AIX NIM Server	This can be used to provide the images during a new AIX installation.
IBM Systems Director HMC extension	This is used to manage the servers and hardware resources that are under the HMC management. It can create or delete new LPARs, and perform power control.

For more information about the IBM Director configuration, features, and extensions, see the following sources:

- IBM Director public library

  http://publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo_all/diricinfoparent.html
- *Implementing IBM Director*, SG24-6188

3.6 Installation tools

In multitenant environment, different tenants require different versions of the OS. Tenants also require fast installation of the complete system. To meet these requirements, we must have automation tools and a well-defined job description.

We have several tools that can support the installation step. In this section, we discuss the various tools and their usage.

3.6.1 Installation with Network Installation Management

Network Installation Management (NIM) environment is an AIX feature that can install AIX OS and packages. It also supports customization during or after the installation. The customization is covered in 3.7.1, “Customization with Network Installation Management” on page 131.

The NIM environment consists of a server and one or multiple clients. It can use any available network, either a physical or virtual adapters, to perform NIM operations using rsh or NIM Service Handler (nimsh) remote commands. The IP addresses have to be predefined before the NIM customization process starts.

The server is called NIM master and the clients are called NIM clients.

Figure 3-40 shows basic NIM resources and when each resource is used.

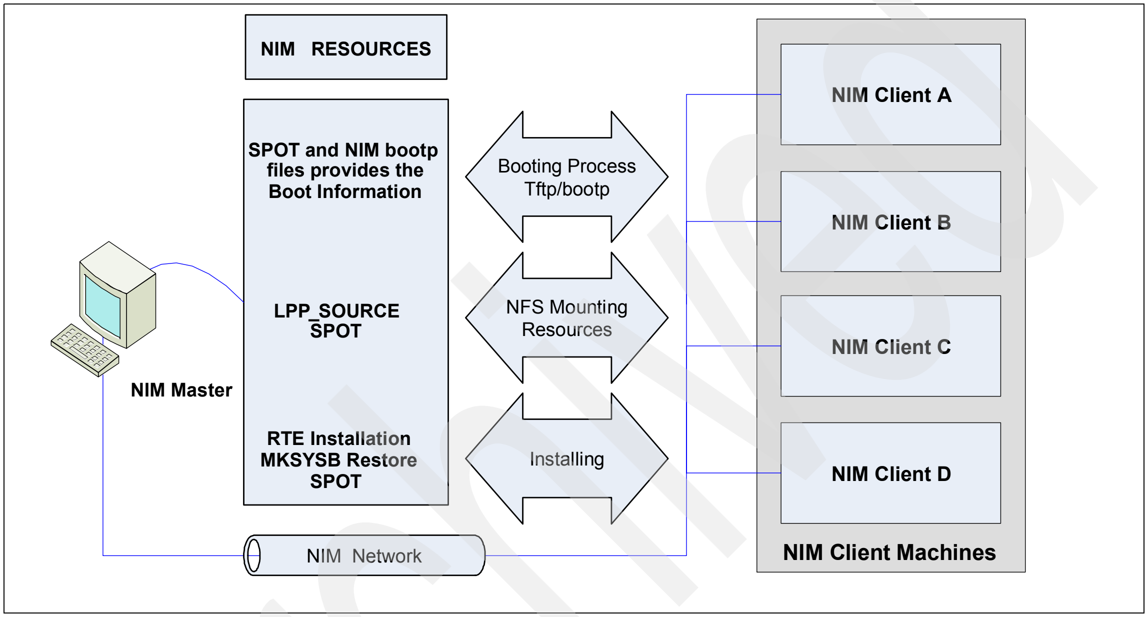

Figure 3-40 Basic NIM resources during the installation process

NIM master and resources

This section introduces the NIM master concepts and the major resources that are used in the NIM environment.

NIM master

The NIM master can be any available AIX machine that uses a physical or a virtual Ethernet adapter. The only requirement is a fully operational network between the NIM master and the NIM clients.

The NIM master can be enabled by installing the `bos.sysmgt.nim.master` file set in the machine that is going to be used as the NIM master.

To configure the NIM environment we can use the **`smit nim`** command (Advanced or Easy Startup options), the nim_master_setup script (available under the `/usr/sbin` directory), or the NIM command line options.

NIM resources

By using NIM, the entire installation process occurs over the network. No AIX CDs are used in the boot or installation phases. For this reason, the NIM requires certain resources to provide all the information necessary during installation.

For basic installation, NIM requires only the Shared Product Object Tree (SPOT) and the licensed program product (LPP) source (lpp_source) resources. If you use these two resources, the NIM requires manual input of information during the installation. The main resources are listed in Table 3-4.

Table 3-4 Main NIM resources and their functions

Resource	**Description**	**Usage**
SPOT	This resource provides /usr information during the installation process. Contains libraries, binary objects, and executables. Replaces the ramdisk available on the installation media.	Used during the boot process to provide binary, library and kernel code. Resources are mounted over the NIM client during the boot phase or are transferred by bootpd and tftpd processes. It can also be used for installation.
lpp_source	This software repository contains a copy of the file sets from the installation media, such as AIX media.	Used during installation and migration. Provides the file set listing that will be copied over the client during the network installation
mksysb	This resource contains the image of the root volume group previously created with the AIX mksysb command.	Can be used to install multiple machines that have the same configuration (such as OS levels and packages)
bosinst_data	This resource is similar to the AIX bosinst.data file. It contains information about console, language, target disks and so on.	Used during the installation of a mksysb resource to provide further information about how the mksysb will be restored
fb_script	This is a predefined script to perform additional customization after the installation’s first boot.	Used only during the first boot after a new installation
script	This resource points to executable and regular files containing executable tasks, usually a shell script.	Used after installation to perform additional customization as file system creation, resizing or to add extra users and groups

> **Note:** After the NIM resources are created, they need to be allocated to the respective client in order to enable the installation.

NIM client

The NIM client can be any supported AIX machine that is able to communicate with the NIM master.

The NIM client can be defined from the NIM master, while the client machine has no OS installed. If we have an existing AIX machine that has to be managed by an existing NIM environment, we can also define it as a NIM client. To configure any existing AIX machine as a NIM client, the requirement, besides the network, is to have the `bos.sysmgt.nim.client` file set, which is installed by default.

> **Reminder:** Any NIM machine, either the master or clients, must be resolved either by DNS or local `/etc/hosts` file. NIM master also requires the CD or DVD allocation to create the first SPOT and lpp_sources, if using standard methods.

To illustrate how to use NIM to customize a new client machine, we created a lab environment by using an IBM Power 570 machine.

Setting up the lab environment

For the lab environment, the configuration shown in Table 3-5 was used.

Table 3-5 Configuration used in the lab environment

Function	Host name	IP address
NIM master	nim1ae0	10.1.1.5
NIM client	nimc1ae0	10.1.1.13
Additional IP for NIM client’s external network	nimc1a	9.3.5.13

> **Note:** All configuration tasks were performed in the NIM master machine. No command or intervention is required by the NIM client during the configuration.

Configuring the NIM master

For this environment we used the Easy Startup option. Figure 3-41 on page 115 is the result of entering the **`smit nim_config_env`** command.

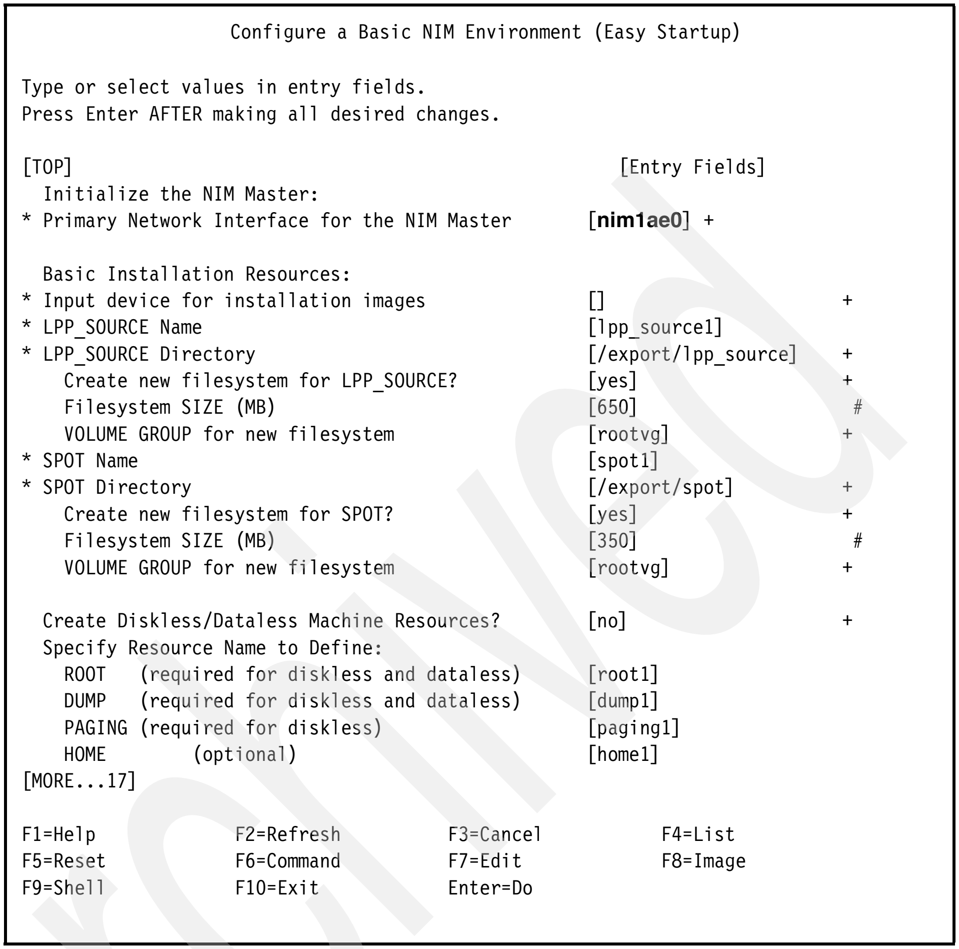

```
                    Configure a Basic NIM Environment (Easy Startup)

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[TOP]                                                   [Entry Fields]
  Initialize the NIM Master:
* Primary Network Interface for the NIM Master       [nim1ae0] +

  Basic Installation Resources:
* Input device for installation images              []                      +
* LPP_SOURCE Name                                   [lpp_source1]
* LPP_SOURCE Directory                              [/export/lpp_source]    +
    Create new filesystem for LPP_SOURCE?           [yes]                   +
    Filesystem SIZE (MB)                            [650]                    #
    VOLUME GROUP for new filesystem                 [rootvg]                +
* SPOT Name                                         [spot1]
* SPOT Directory                                    [/export/spot]          +
    Create new filesystem for SPOT?                 [yes]                   +
    Filesystem SIZE (MB)                            [350]                    #
    VOLUME GROUP for new filesystem                 [rootvg]                +

  Create Diskless/Dataless Machine Resources?       [no]                    +
  Specify Resource Name to Define:
    ROOT   (required for diskless and dataless)     [root1]
    DUMP   (required for diskless and dataless)     [dump1]
    PAGING (required for diskless)                  [paging1]
    HOME        (optional)                          [home1]
[MORE...17]

F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 3-41 Output of the smit nim_config_env command

In this step, we have to insert the host name (`nim1ae0`), which was previously added to the hosts file. Accept the default options and press Enter.

The NIM master is configured with predefined values. Now we can define the NIM client.

Note: Usually a separated network is defined only to be used by the NIM environment, but NIM can use any available network.

Configuring the NIM client

During the NIM master definition, the first SPOT and LPP_SOURCEs were created. To use a machine as the NIM client we have to define it, as described in this section.

Use nim_mkmac fastpath, as shown in Figure 3-42.

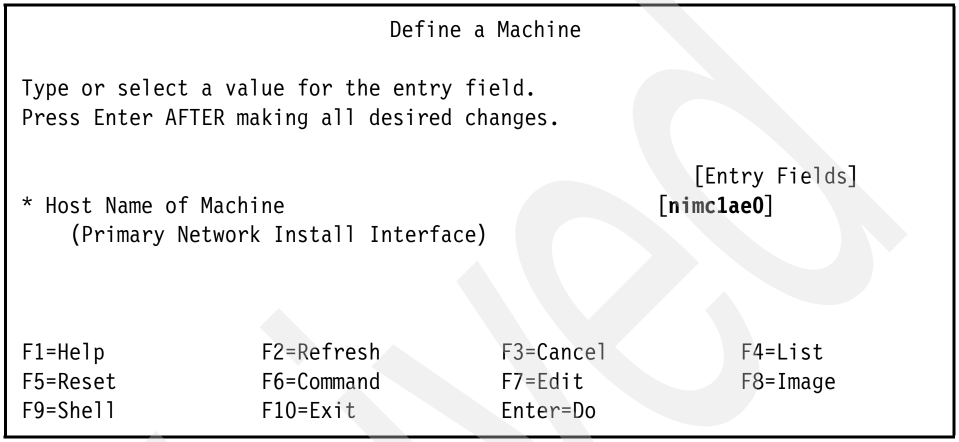

```
                               Define a Machine

Type or select a value for the entry field.
Press Enter AFTER making all desired changes.

                                                             [Entry Fields]
* Host Name of Machine                                    [nimc1ae0]
    (Primary Network Install Interface)



F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 3-42 Output of the smit nim_mkmac command

In this step the `Host Name of Machine` is the resolvable host name. After entering the host name, the panel shown in Figure 3-43 on page 117 is displayed.

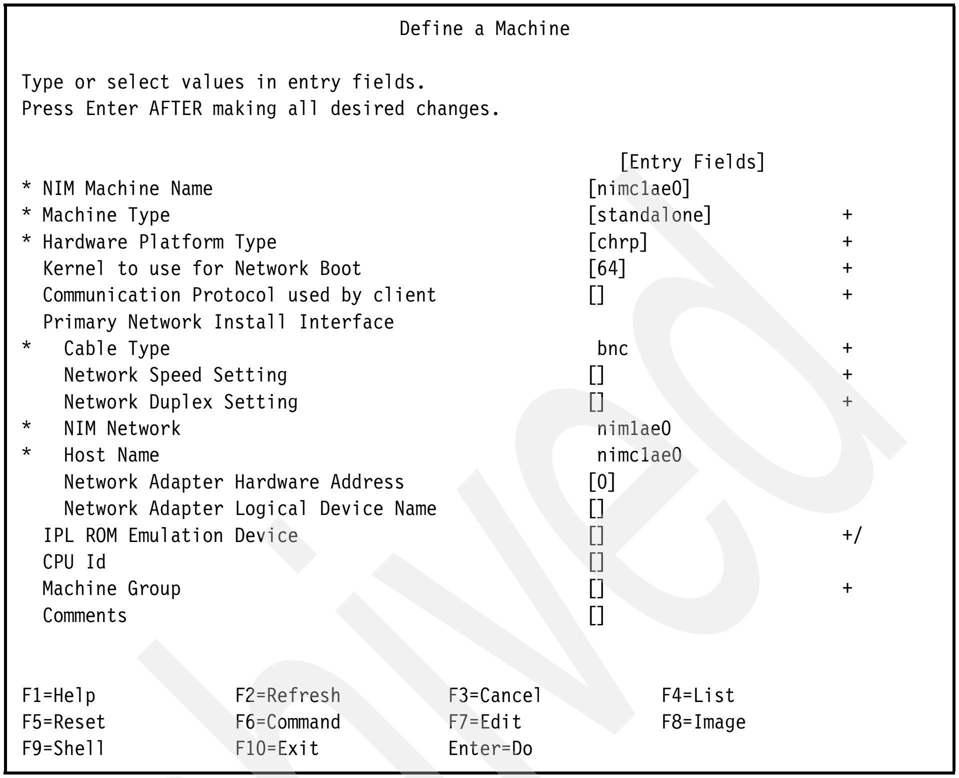

```
                                   Define a Machine

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                                           [Entry Fields]
* NIM Machine Name                                      [nimc1ae0]
* Machine Type                                          [standalone]           +
* Hardware Platform Type                                [chrp]                 +
  Kernel to use for Network Boot                        [64]                   +
  Communication Protocol used by client                 []                     +
  Primary Network Install Interface
*   Cable Type                                           bnc                   +
    Network Speed Setting                               []                     +
    Network Duplex Setting                              []                     +
*   NIM Network                                          nim1ae0
*   Host Name                                            nimc1ae0
    Network Adapter Hardware Address                    [0]
    Network Adapter Logical Device Name                 []
  IPL ROM Emulation Device                              []                     +/
  CPU Id                                                []
  Machine Group                                         []                     +
  Comments                                              []


F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 3-43   Continuation of output of the smit nim_mkmac command

If the host name is not defined either locally or in the DNS server, the panel shown in Figure 3-44 on page 118 is displayed.

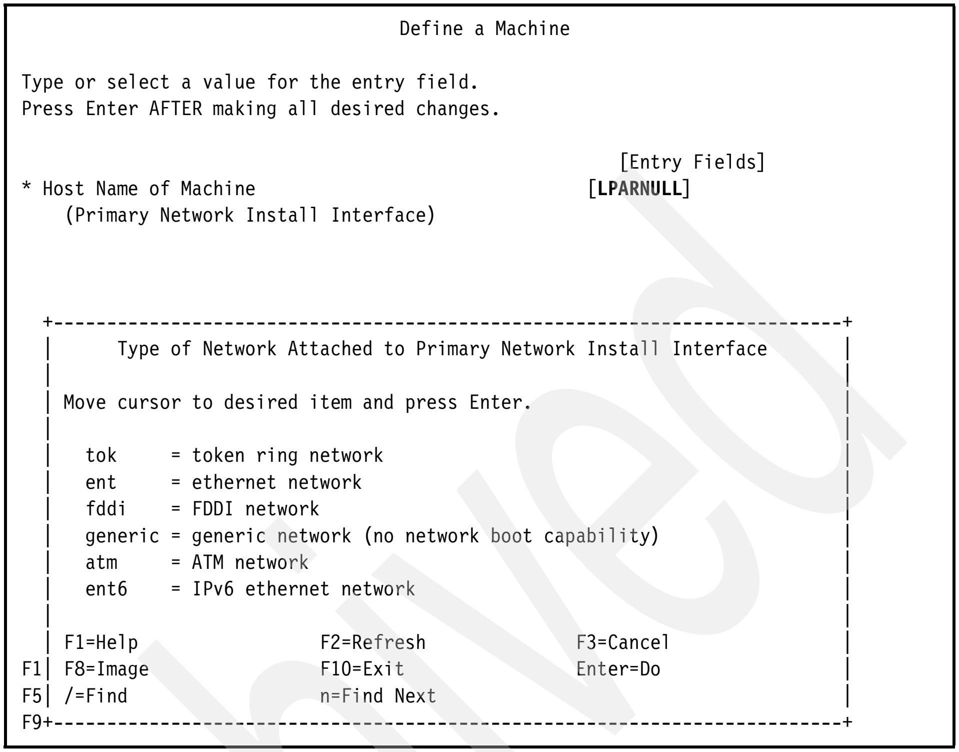

```
                                  Define a Machine

Type or select a value for the entry field.
Press Enter AFTER making all desired changes.

                                                        [Entry Fields]
* Host Name of Machine                               [LPARNULL]
    (Primary Network Install Interface)




  +--------------------------------------------------------------------------+
  |      Type of Network Attached to Primary Network Install Interface       |
  |                                                                          |
  | Move cursor to desired item and press Enter.                             |
  |                                                                          |
  |   tok     = token ring network                                           |
  |   ent     = ethernet network                                             |
  |   fddi    = FDDI network                                                  |
  |   generic = generic network (no network boot capability)                 |
  |   atm     = ATM network                                                  |
  |   ent6    = IPv6 ethernet network                                        |
  |                                                                          |
  | F1=Help                F2=Refresh              F3=Cancel                 |
F1| F8=Image               F10=Exit                Enter=Do                  |
F5| /=Find                 n=Find Next                                       |
F9+--------------------------------------------------------------------------+
```

Figure 3-44 Panel that displays if a non resolvable host name is used

To correctly define the NIM client machine, add the respective `/etc/hosts` information and repeat the task.

After completing this task the machine is ready to be used (with default definitions).

Understanding the installation methods

NIM provides three installation methods (types):

mksysb This type uses an existing AIX mksysb image and restores it into the client machine.

rte This type uses an existing lpp_source to install a new environment or to migrate an existing system.

spot This type uses a predefined SPOT and copies its contents to the NIM client machine.

In “Starting the system WPAR” on page 106, we use the **`rte`** installation type.

Starting the NIM client installation

To start the basic installation we used the **`smit nim_bosinst`** command, as shown in Figure 3-45.

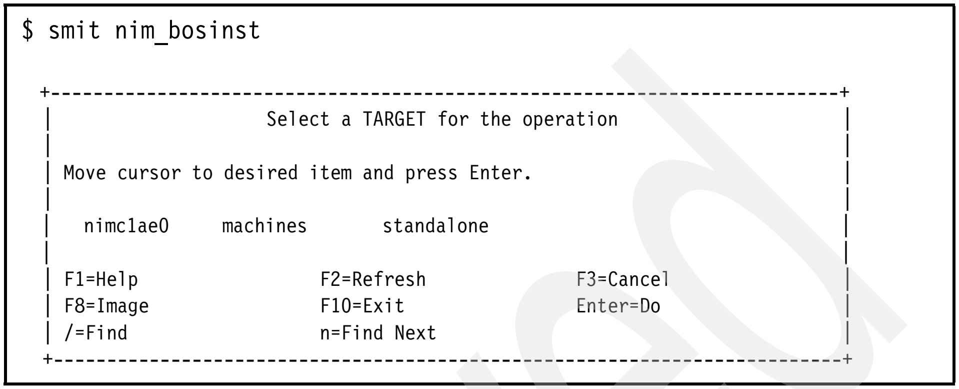

```
$ smit nim_bosinst

  +--------------------------------------------------------------------------+
  |                   Select a TARGET for the operation                      |
  |                                                                          |
  | Move cursor to desired item and press Enter.                             |
  |                                                                          |
  |   nimc1ae0     machines       standalone                                 |
  |                                                                          |
  | F1=Help                 F2=Refresh              F3=Cancel                |
  | F8=Image                F10=Exit                Enter=Do                 |
  | /=Find                  n=Find Next                                      |
  +--------------------------------------------------------------------------+
```

Figure 3-45 Output of the smit nim_bosinst command

Select the NIM client machine. The installation types are listed, as shown in Figure 3-46.

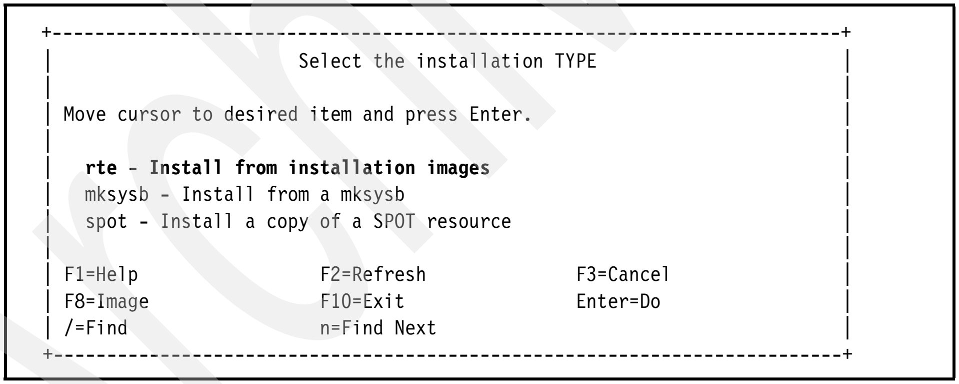

```
  +--------------------------------------------------------------------------+
  |                      Select the installation TYPE                        |
  |                                                                          |
  | Move cursor to desired item and press Enter.                             |
  |                                                                          |
  |   rte - Install from installation images                                 |
  |   mksysb - Install from a mksysb                                         |
  |   spot - Install a copy of a SPOT resource                               |
  |                                                                          |
  | F1=Help                 F2=Refresh              F3=Cancel                |
  | F8=Image                F10=Exit                Enter=Do                 |
  | /=Find                  n=Find Next                                      |
  +--------------------------------------------------------------------------+
```

Figure 3-46 Installation types panel

For this installation the `rte` type was selected. The LPP_SOURCE selections are listed, as shown in Figure 3-47 on page 120.

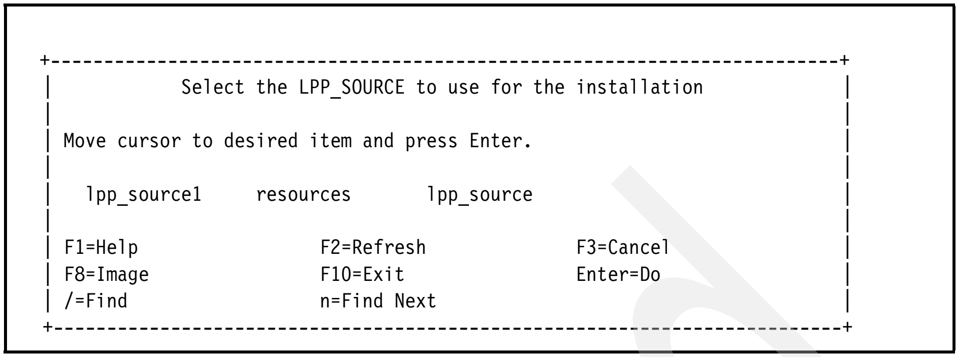

```
+--------------------------------------------------------------------------+
|           Select the LPP_SOURCE to use for the installation              |
|                                                                          |
| Move cursor to desired item and press Enter.                             |
|                                                                          |
|   lpp_source1     resources       lpp_source                             |
|                                                                          |
| F1=Help                F2=Refresh             F3=Cancel                  |
| F8=Image               F10=Exit               Enter=Do                   |
| /=Find                 n=Find Next                                       |
+--------------------------------------------------------------------------+
```

Figure 3-47 LPP_Source selection

After selecting the LPP_SOURCE, we select the SPOT to be used, as shown in Figure 3-48.

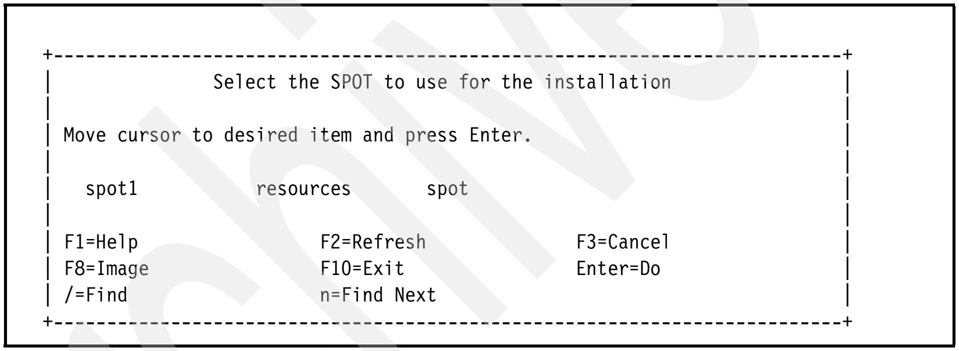

```
+--------------------------------------------------------------------------+
|              Select the SPOT to use for the installation                 |
|                                                                          |
| Move cursor to desired item and press Enter.                             |
|                                                                          |
|   spot1           resources       spot                                   |
|                                                                          |
| F1=Help                F2=Refresh             F3=Cancel                  |
| F8=Image               F10=Exit               Enter=Do                   |
| /=Find                 n=Find Next                                       |
+--------------------------------------------------------------------------+
```

Figure 3-48 SPOT selection

After selecting the SPOT, other installation options are listed, as shown in Figure 3-49 on page 121.

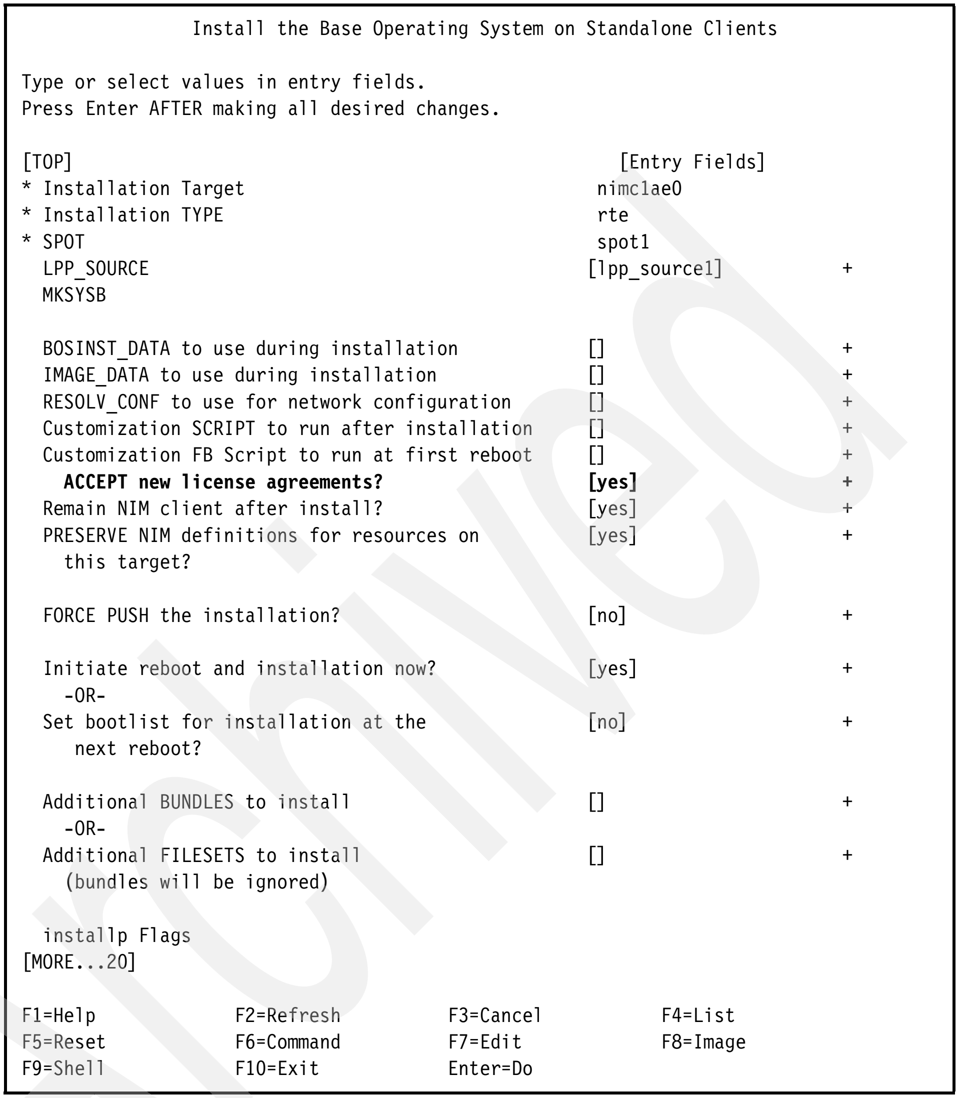

```
                Install the Base Operating System on Standalone Clients

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[TOP]                                                   [Entry Fields]
* Installation Target                                 nimc1ae0
* Installation TYPE                                   rte
* SPOT                                                spot1
  LPP_SOURCE                                         [lpp_source1]           +
  MKSYSB

  BOSINST_DATA to use during installation            []                      +
  IMAGE_DATA to use during installation              []                      +
  RESOLV_CONF to use for network configuration       []                      +
  Customization SCRIPT to run after installation     []                      +
  Customization FB Script to run at first reboot     []                      +
    ACCEPT new license agreements?                   [yes]                   +
  Remain NIM client after install?                   [yes]                   +
  PRESERVE NIM definitions for resources on          [yes]                   +
    this target?

  FORCE PUSH the installation?                       [no]                    +

  Initiate reboot and installation now?              [yes]                   +
    -OR-
  Set bootlist for installation at the               [no]                    +
     next reboot?

  Additional BUNDLES to install                      []                      +
    -OR-
  Additional FILESETS to install                     []                      +
    (bundles will be ignored)

  installp Flags
[MORE...20]

F1=Help             F2=Refresh          F3=Cancel           F4=List
F5=Reset            F6=Command          F7=Edit             F8=Image
F9=Shell            F10=Exit            Enter=Do
```

Figure 3-49 Default options for the installation menu

In the menu, change the `ACCEPT new license agreements` question to `yes`, in order to automatically accept any license requirement during installation.

The default installation process automatically begins.

Note: For this example the default installation (SPOT and LPP_SOURCE) was selected. We did not show all the installation windows.

For more details about configuring the NIM environment to provide further customization, see 3.7.1, “Customization with Network Installation Management” on page 131

For more information about NIM configuration and installation, see *NIM from A to Z in AIX 5L*, SG24-7296.

3.6.2 Installation with Tivoli Provisioning Manager

In 3.5.3, “Allocation with Tivoli Provisioning Manager” on page 80, we created LPARs on the host platform. Now we can install an AIX OS image on the virtual sever.

We can use ADS, Jumpstart, Kickstart, NIM, Rembo (recently renamed Tivoli Provisioning Manager for OS Deployment), and YaST as a TPM boot server. For managing AIX OS, NIM is recommended. We define a NIM server as a TPM boot server.

Note: When creating NIM server, be sure to create the mksysb image on the NIM server. When TPM discovers the NIM, it requires a mksysb image as an install image.

To begin the installation:

1. Select **Inventory** → **Infrastructure Management** → **Boot Servers**.
2. In the **Manage Boot Servers** window, select **Edit** → **Add Boot Server Wizard**, as shown in Figure 3-50.

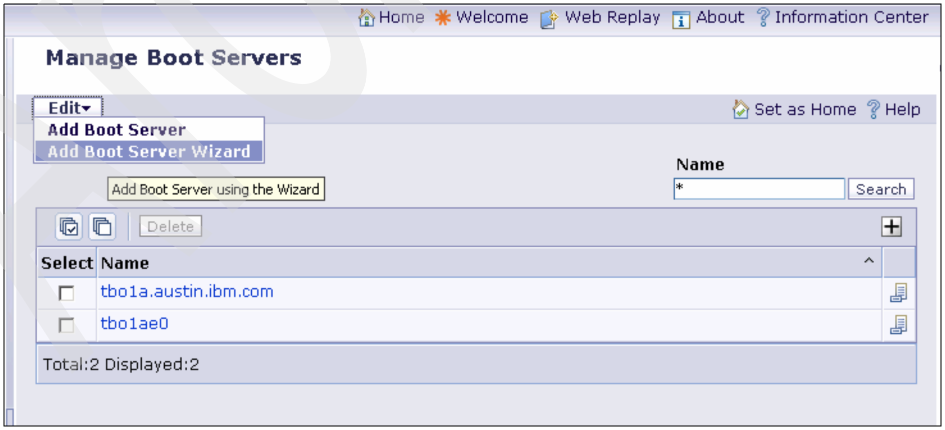

Name
*

Select	Name
□	tbo1a.austin.ibm.com
□	tbo1ae0

Total:2 Displayed:2

Figure 3-50 Create boot server in TPM with Add Boot Server Wizard

3. Select **NIM** as the boot server type, as shown in Figure 3-51.

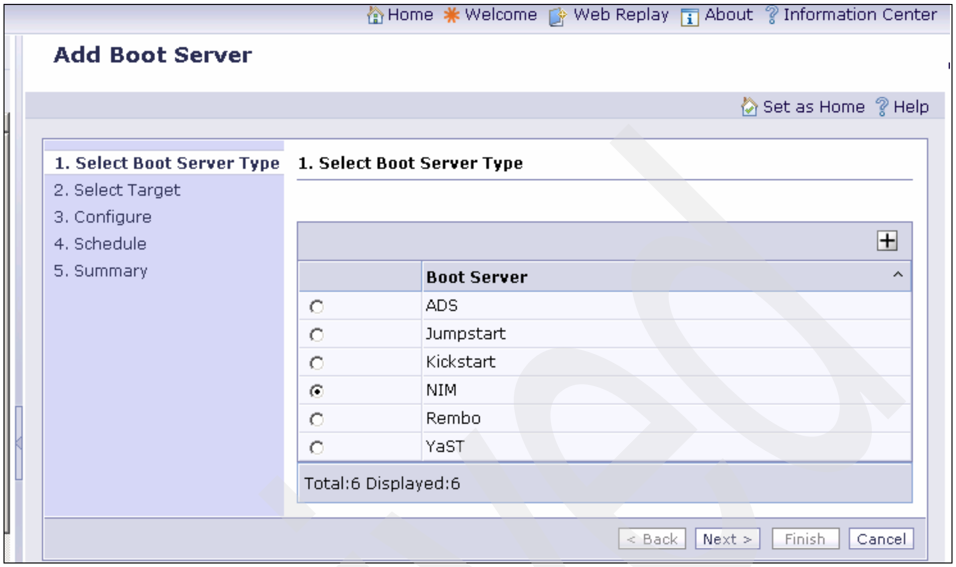

Add Boot Server

1. Select Boot Server Type
2. Select Target
3. Configure
4. Schedule
5. Summary

1. Select Boot Server Type

	Boot Server
○	ADS
○	Jumpstart
○	Kickstart
●	NIM
○	Rembo
○	YaST

Total:6 Displayed:6

Figure 3-51 Select Boot Server Type

4. Select NIM server as the target server, as shown in Figure 3-52 on page 124.

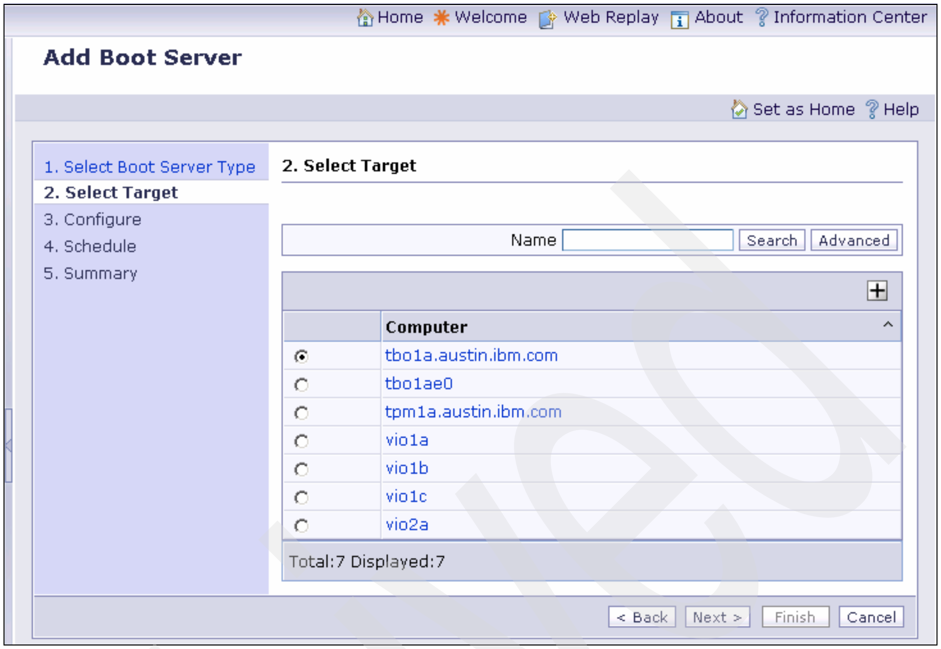

Add Boot Server

2. Select Target

Name

	Computer
●	tbo1a.austin.ibm.com
○	tbo1ae0
○	tpm1a.austin.ibm.com
○	vio1a
○	vio1b
○	vio1c
○	vio2a

Total:7 Displayed:7

Figure 3-52 Select NIM server in the Computer section

Figure 3-53 shows that we registered NIM server as a boot server for TPM.

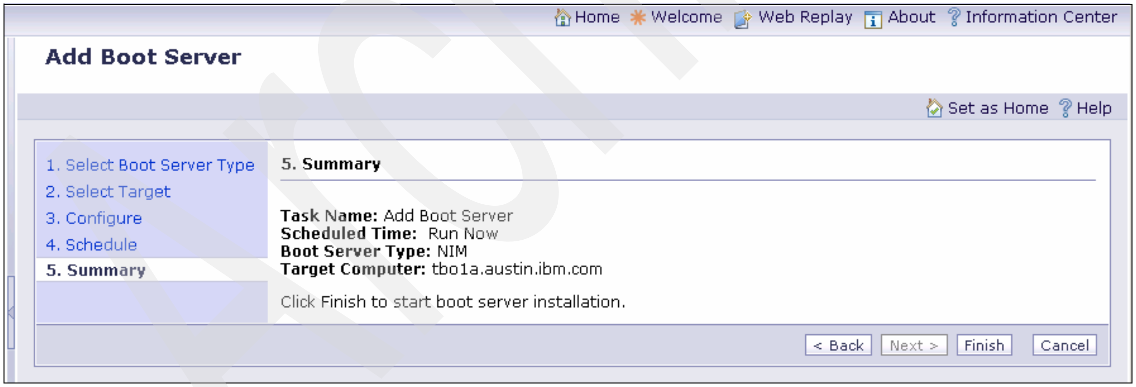

Add Boot Server

5. Summary

Task Name: Add Boot Server
Scheduled Time: Run Now
Boot Server Type: NIM
Target Computer: tbo1a.austin.ibm.com

Click Finish to start boot server installation.

Figure 3-53 Define NIM server as a Boot server in TPM

5. Click **Finish**.
6. Select **Software Management** → **Manage Software Catalog** → **Images**. Figure 3-54 on page 125 shows the mksysb image on the NIM server.

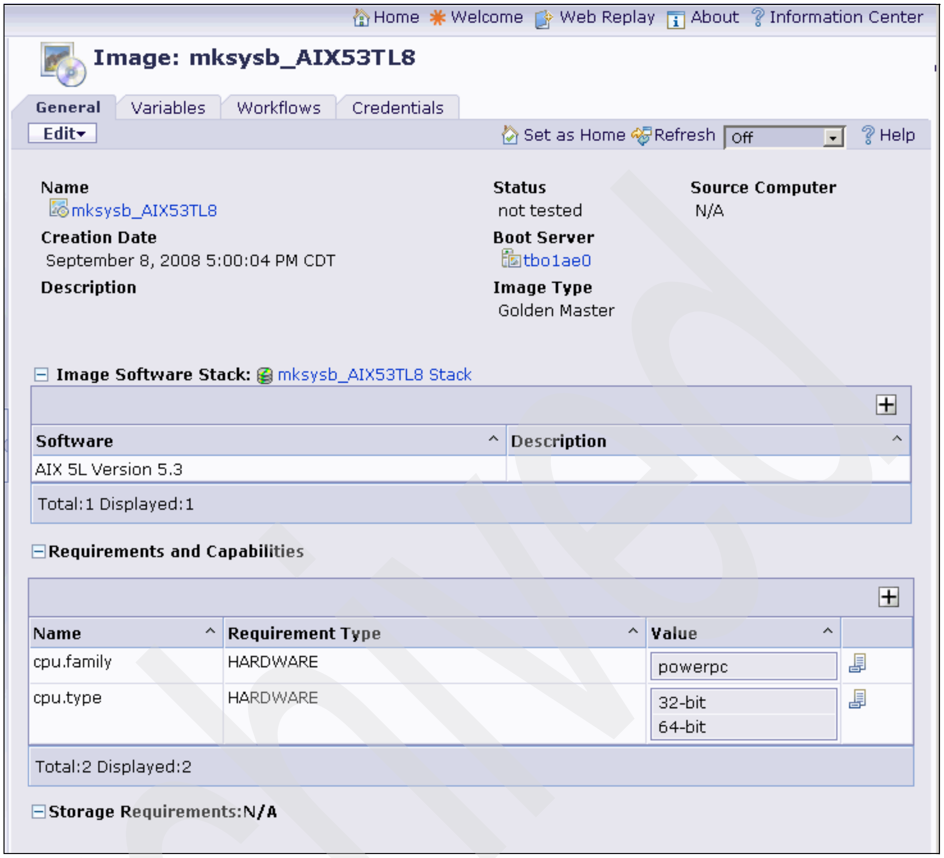

Image: mksysb_AIX53TL8

Name
mksysb_AIX53TL8

Creation Date
September 8, 2008 5:00:04 PM CDT

Description

Status
not tested

Boot Server
tbo1ae0

Image Type
Golden Master

Source Computer
N/A

Image Software Stack: mksysb_AIX53TL8 Stack

Software	Description
AIX 5L Version 5.3	

Total:1 Displayed:1

Requirements and Capabilities

Name	Requirement Type	Value
cpu.family	HARDWARE	powerpc
cpu.type	HARDWARE	32-bit 64-bit

Total:2 Displayed:2

Storage Requirements:N/A

Figure 3-54 The mksysb image of NIM discovered by TPM

Before we can install the OS image to specify the LPAR, we must first create the NIM server client in the NIM server we registered. It is enough to add the host name to `/etc/hosts` file in NIM server. The work flows are shown in Figure 3-55 on page 126.

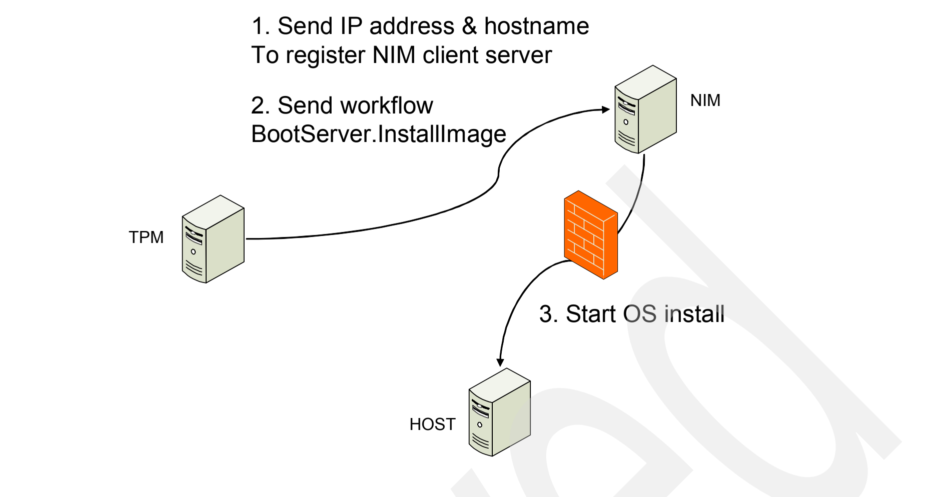

Figure 3-55   Registering NIM client in NIM server to install OS

We used the script shown in Example 3-11 to register a NIM client.

Example 3-11   Shell script that registers the host name and IP address in NIM

```
#!/usr/bin/ksh

NIM_SERVER=tbo1a.austin.ibm.com
TMP_FILE=/tmp/tmp.data

if [ $1 ]
then
        rsh $NIM_SERVER -l root cat /etc/hosts > $TMP_FILE
        echo "10.1.1.$2  $1" >> $TMP_FILE
        rcp $TMP_FILE $NIM_SERVER:/etc/hosts
        rm $TMP_FILE
else
        echo "Usage:nim_register [hostname] [ip address]"
        echo "Example) nim_register sarah 17"
fi
```

To complete the installation:

1. Select **Software Management** → **Manage Software Catalog** → **Images**.
2. Click on the image to install and select **Edit** → **Install Image**.
3. Select the computer on which to install the OS, as shown in Figure 3-56.

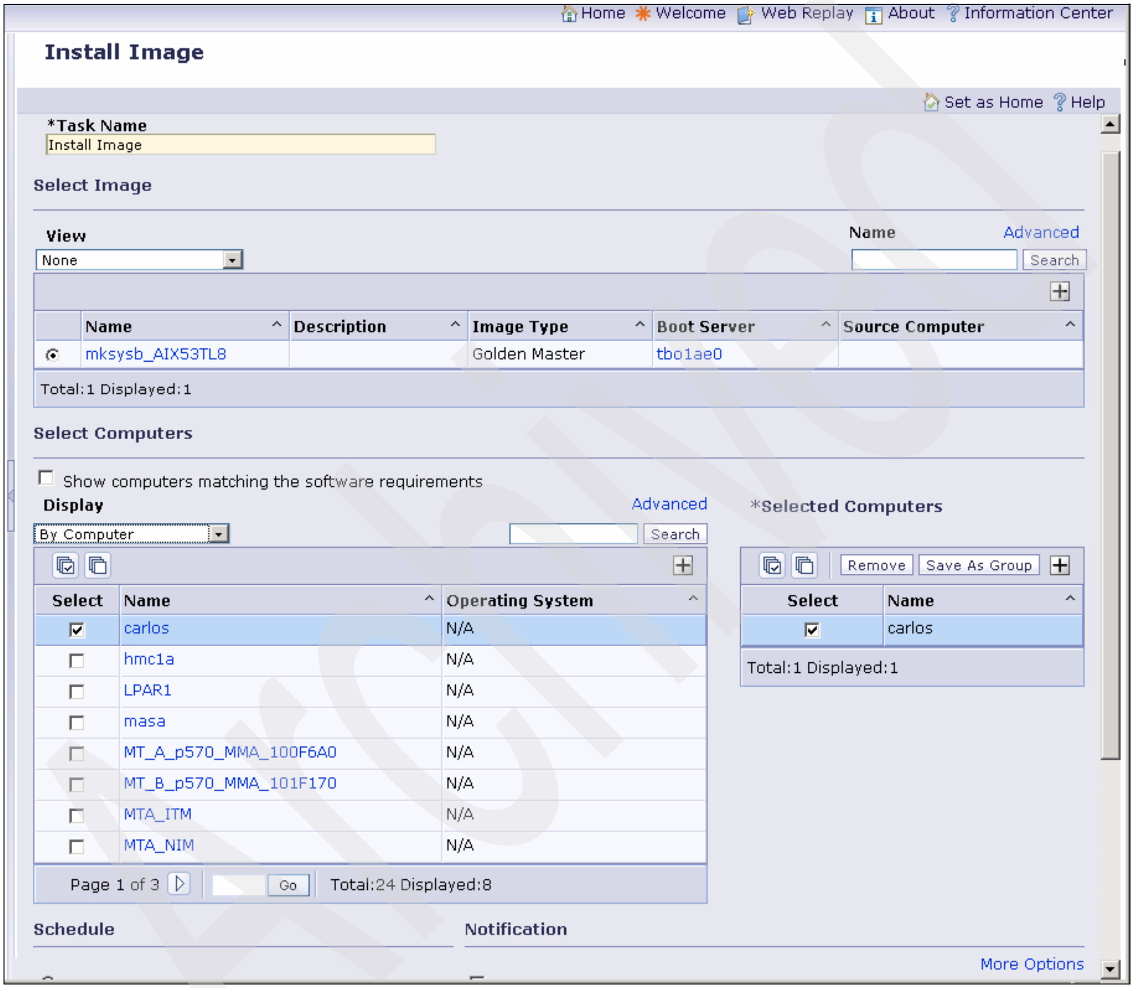

Install Image

*Task Name: Install Image

Select Image

View: None

Name: Search

	Name	Description	Image Type	Boot Server	Source Computer
●	mksysb_AIX53TL8		Golden Master	tbo1ae0	

Total:1 Displayed:1

Select Computers

☐ Show computers matching the software requirements

Display: By Computer

Select	Name	Operating System
☑	carlos	N/A
☐	hmc1a	N/A
☐	LPAR1	N/A
☐	masa	N/A
☐	MT_A_p570_MMA_100F6A0	N/A
☐	MT_B_p570_MMA_101F170	N/A
☐	MTA_ITM	N/A
☐	MTA_NIM	N/A

Page 1 of 3 Go Total:24 Displayed:8

*Selected Computers

Select	Name
☑	carlos

Total:1 Displayed:1

Schedule

Notification

Figure 3-56 Image Install options

TPM automatically starts the local partition and starts the OS installation. After installation is completed, we have a running partition AIX installed as shown in Figure 3-57 on page 128.

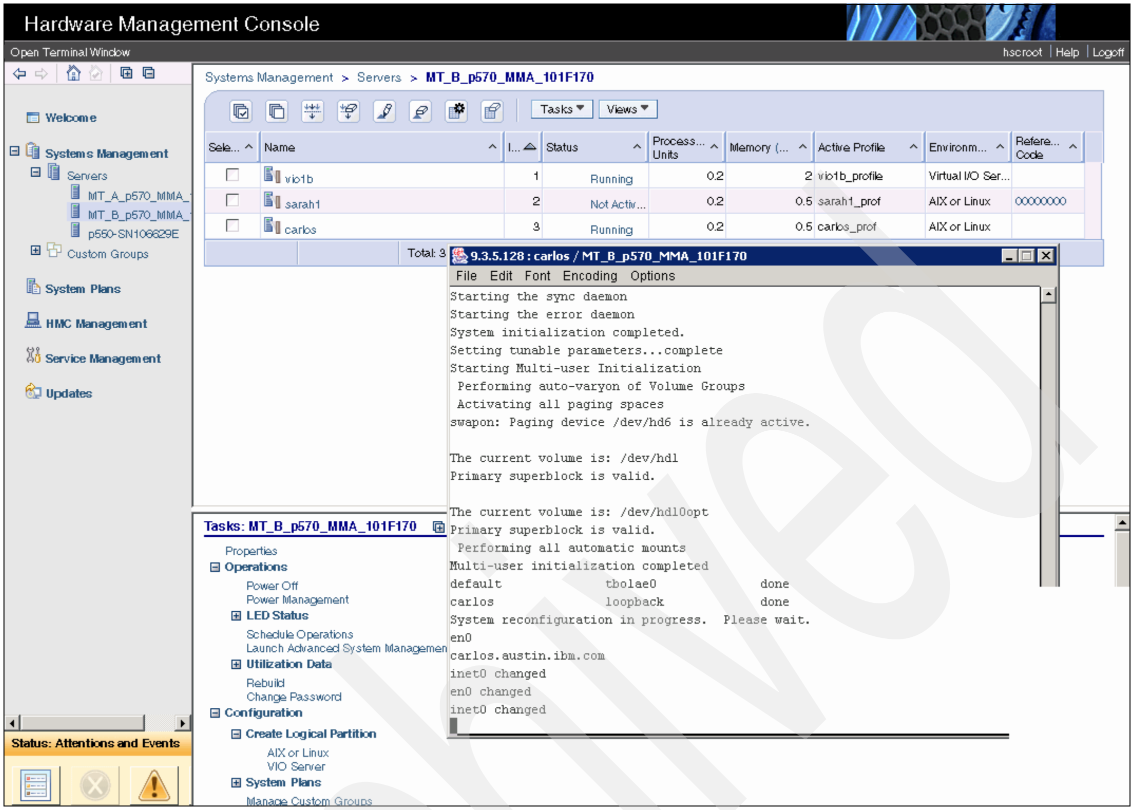

Sele...	Name	I...	Status	Process... Units	Memory (...	Active Profile	Environm...	Refere... Code
□	vio1b	1	Running	0.2	2	vio1b_profile	Virtual I/O Ser...	
□	sarah1	2	Not Activ...	0.2	0.5	sarah1_prof	AIX or Linux	00000000
□	carlos	3	Running	0.2	0.5	carlos_prof	AIX or Linux	

```
Starting the sync daemon
Starting the error daemon
System initialization completed.
Setting tunable parameters...complete
Starting Multi-user Initialization
 Performing auto-varyon of Volume Groups
 Activating all paging spaces
swapon: Paging device /dev/hd6 is already active.

The current volume is: /dev/hd1
Primary superblock is valid.

The current volume is: /dev/hd10opt
Primary superblock is valid.
 Performing all automatic mounts
Multi-user initialization completed
default             tbolae0                done
carlos              loopback               done
System reconfiguration in progress.  Please wait.
en0
carlos.austin.ibm.com
inet0 changed
en0 changed
inet0 changed
```

Figure 3-57 Running the OS after installation

Web Replay

To automate the most common tasks with the TPM user interface, TPM provides an automation function called Web Replay.

Web Replay enables you to capture the steps you go through when solving a task. The steps are saved in a *scenario*. The scenario can be reused by other users to perform the same task with a single click.

With Web Replay you can automate complex tasks such as discovery, virtual server creation, software installation, OS management, and patch management. You can record the mouse actions required to implement the task and play back your recording later, as necessary.

The Web Replay function and manuals are accessed from the Welcome page of TPM, as shown in Figure 3-58 on page 129.

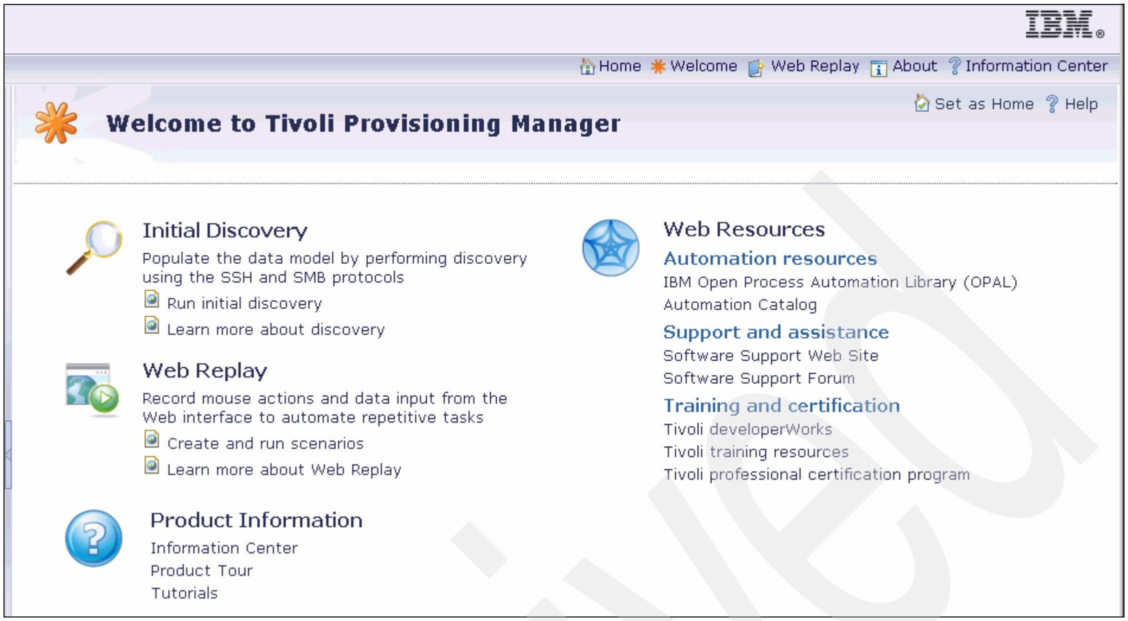

IBM.

Home Welcome Web Replay About Information Center

Set as Home Help

Welcome to Tivoli Provisioning Manager

Initial Discovery

Populate the data model by performing discovery using the SSH and SMB protocols

Run initial discovery

Learn more about discovery

Web Replay

Record mouse actions and data input from the Web interface to automate repetitive tasks

Create and run scenarios

Learn more about Web Replay

Product Information

Information Center

Product Tour

Tutorials

Web Resources

Automation resources

IBM Open Process Automation Library (OPAL)

Automation Catalog

Support and assistance

Software Support Web Site

Software Support Forum

Training and certification

Tivoli developerWorks

Tivoli training resources

Tivoli professional certification program

Figure 3-58 Welcome page of Tivoli Provisioning Manager

To learn how to use Web Replay, download instructional demonstrations from:

http://demos.dfw.ibm.com

To use Web Replay for specific purposes, download best practices from:

http://www.ibm.com/software/brandcatalog/portal/opal

3.6.3 Installation with IBM Director Virtual Image Manager

IBM Director Virtual Image Manager is part of the Virtualization Manager extension to IBM Director. It provides support for creating and managing images to install OS on virtual servers. IBM Virtual Image Manager supports deployment under the control of HMC or IVM and has the following features:

- System templates management
- Virtual server image management with OS image
- Create and clone virtual server with OS image
- Install OS

For more information about IBM Virtual Image Manager, visit:

http://www.ibm.com/systems/management/director/extensions/vim.html

3.7 Customization tools

In a multitenant environment, we can provide physical or virtual machines, install the appropriate operating system, and perform additional customization during the deployment of the new resource.

Each tenant has its own network, environment, and operating-system-level requirements. Certain settings can be changed during installation time, either directly after the first boot or in a maintenance-scheduled operation at a later time, as shown in Figure 3-59.

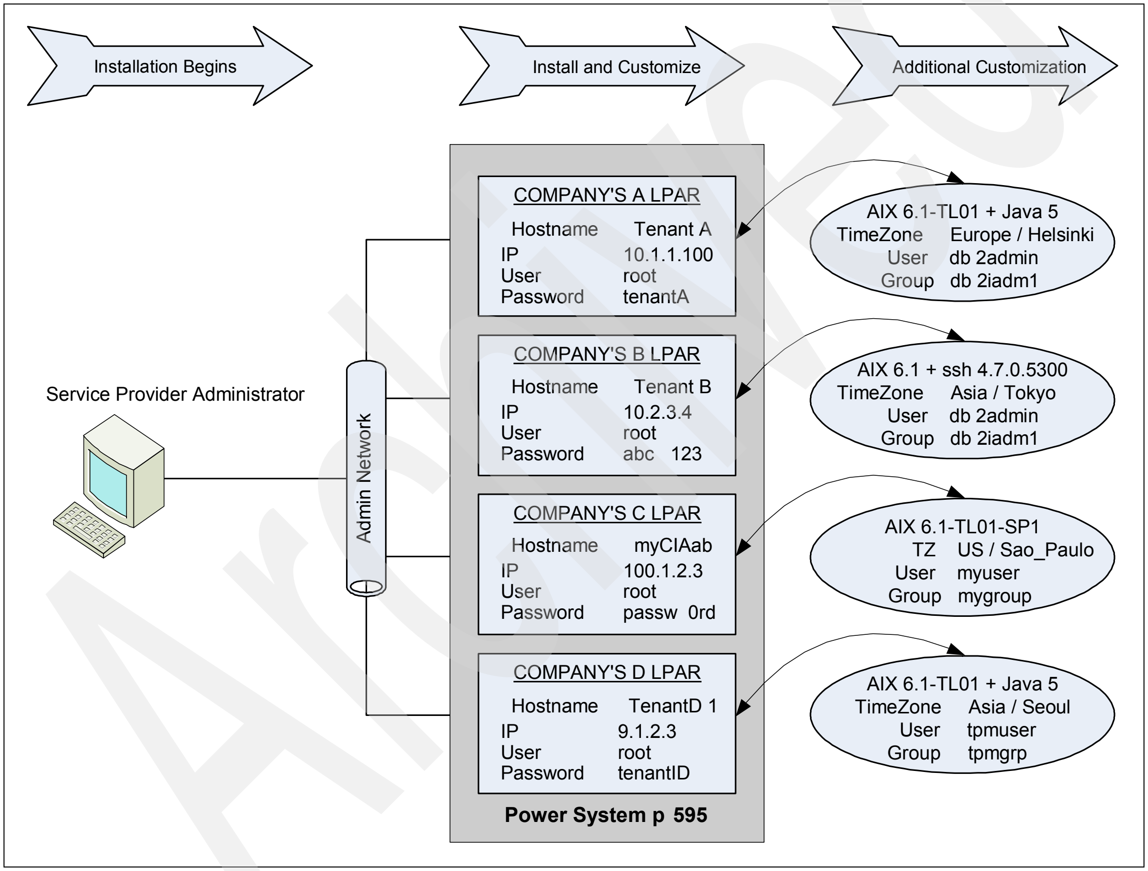

Figure 3-59 Graphical flow for some customizable requirements

Several ways are available for performing automated customization in a new or current managed environment. We can customize the system by using shell scripts, and AIX functions and tools. In this book, we use NIM and the TPM.

In this section we show how to perform and automate customization tasks in our managed environment by using the various automation tools.

3.7.1 Customization with Network Installation Management

As discussed in 3.6.1, "Installation with Network Installation Management" on page 111, the NIM uses certain resources to install the basic AIX file sets or install a customized mksysb image.

NIM can also perform customization tasks during or after installation of the OS.

Table 3-6 lists several NIM resources that can be used to provide customization and when they can be used.

Table 3-6 Resource types and descriptions

Type of resource	Description	Usage
lpp_source	This software repository is for additional AIX packages and rpm file sets.	Used during installation, migration. It can also be used to update the AIX of an existing machine or to install new file sets and bundles.
bosinst_data	This resource is similar to the AIX bosinst.data file and contains information about console, language, target disks, and so on.	Used during the installation time of a mksysb resource to provide additional information about how the mksysb will be restored and to avoid manual interaction
fb_script	This resource is a predefined script to perform additional customization after the installation's first boot.	Used only during the first boot after a new installation to perform the first customization tasks
script	This resource points to executable and regular file containing executable tasks, usually a shell script.	Used after the installation to perform additional customization as file system creation, resizing, or to add extra users and groups

NIM resources for customization

The following NIM resources can be used to perform various configuration tasks:

- Perform a base operating system (BOS) installation.
- Perform a runtime execution (RTE) installation.
- Create a mksysb resource to duplicate the same customized installation image between multiple machines.
- Create several SPOTs and LPP_SOURCES with different levels of AIX and technology levels.
- Create a group of NIM client machines, containing all machines that are under the same tenant policies and run customized tasks for the entire group.
- Create customized scripts to run during or after the installation.

Customization during installation

In this section, we show how to create and use several NIM resources to perform customization during BOS, RTE, or mksysb installation.

The following resources can be used separately during any NIM installation:

- NIM SPOT and lpp_source resources
- NIM bosinst_data
- NIM fb_script resource

For practical purposes we first defined all the resources and performed a single installation to show the results.

Creating NIM SPOT and lpp_source

The modified SPOT or lpp_source can be used to install a specific version of OS. It can also contain additional file set bundles to support the tenant's applications.

To create a new SPOT or LPP_Source, AIX CDs must be copied into any target file system. For this example, the AIXs files were copied (bffcreate format) into the `/software/aix-6100-01` file system and we created a new lpp_source containing full AIX installation CDs set.

Listing previous lpp_sources by using the **`lsnim`** shown in Example 3-12.

Example 3-12 Listing lpp_sources

```
$ lsnim -t lpp_source
lpp_aix61TL1          resources       lpp_source
lpp_aix53TL8          resources       lpp_source
```

Defining a new lpp_source called `lpp_aix61TL1_Full` is shown in Example 3-13.

Example 3-13   Output generated by the nim -o define command

```
$ nim -o define -t lpp_source -a server=master -a
location=/software/aix-6100-01-01/installp/ppc lpp_aix61TL1_Full

Preparing to copy install images (this will take several minutes)...
Now checking for missing install images...
All required install images have been found. This lpp_source is now ready.
```

Listing the created lpp_source and its contents by using **`lsnim`** command is shown in Example 3-14.

Example 3-14   Output of lsnim and nim -o showres command

```
$ lsnim -t lpp_source
lpp_aix61TL1          resources       lpp_source
lpp_aix53TL8          resources       lpp_source
lpp_aix61TL1_Full     resources       lpp_source

$ nim -o showres lpp_aix61TL1_Full
(multiple lines were removed)
#   XL SMP Runtime Messages - Japanese
xlsmp.msg.ZH_CN.rte         1.7.0.0                    I  N usr
#   XL SMP Runtime Messages - Simplified Chinese UTF-8
xlsmp.msg.Zh_CN.rte         1.7.0.0                    I  N usr
#   XL SMP Runtime Messages - Simplified Chinese
xlsmp.msg.en_US.rte         1.7.0.0                    I  N usr
#   XL SMP Runtime Messages - U.S. English
xlsmp.msg.ja_JP.rte         1.7.0.0                    I  N usr
#   XL SMP Runtime Messages - Japanese IBM-eucJP
xlsmp.msg.zh_CN.rte         1.7.0.0                    I  N usr
#   XL SMP Runtime Messages - Simplified Chinese IBM-euc
xlsmp.rte                   1.7.0.0                    I  N usr
#   SMP Runtime Library
```

Creating NIM bosinst_data

To create this resource we use any regular bosinst_data file, modify the file, and create a bosinst_data NIM resource.

The bosinst_data NIM resource can be used to perform some customization during the installation process. For example, set the NIM installation process to avoid manual intervention (PROMPT = no), enable TCB (Trusted Computing Base), or install additional packages.

> **Note:** The AIX provides a default bosinst_data file, named bosinst.template, that resides in the `/usr/lpp/bosinst` directory and can be used as the sample for the NIM bosinst_data resource

For this example, we modified the NIM master's `/var/adm/ras/bosinst_data` file and saved it as `/var/adm/ras/bosinst_data_nimc1ae0`, shown in Example 3-15.

Example 3-15 /var/adm/ras/bosinst_data_nimc1ae0 file

```
(multiple lines removed)
control_flow:
    CONSOLE = Default
    INSTALL_METHOD = overwrite
    PROMPT = no
    EXISTING_SYSTEM_OVERWRITE = yes
    INSTALL_X_IF_ADAPTER = yes
 (removed lines)
    TCB = yes
(removed lines)
    SERVER_BUNDLE = yes
(removed lines)
```

The file was changed to not to show the prompt, to enable TCB, and to install the server bundle packages during the OS installation.

Now, we define a bosinst_data resource to use the created file by using the **`nim -o define`** command, as shown in Example 3-16.

Example 3-16 Defining the bosinst_data resource

```
$ nim -o define -t bosinst_data -a server=master -a
location=/var/adm/ras/bosinst_data_nimc1ae0 bosinsts_data_nimc1ae0
$ lsnim -t bosinst_data
bosinsts_data_nimc1ae0     resources       bosinst_data
```

Creating NIM fb_script

To provide additional customization, we created a NIM resource named fb_script. It changes the TZ variable, the fsize ulimit parameter, and configures the IP address for the additional Ethernet network interface en1 (used to communicate with external network).

A script named `/var/adm/ras/fb_script_nimc1ae0.sh` was created with the contents shown in Example 3-17 on page 135.

Example 3-17 Script /var/adm/ras/fb_script_nimc1ae0.sh

```
# !/bin/ksh
# Script fb_script_nimc1ae0.sh
# This script was created to customize the NIM Client nimc1ae0
#
# Created in 2008
# Objective : Customize TimeZone, ulimit and Network Information
#
# Redirect the STDOUT and STDERR to a logfile
exec 1> /tmp/${0##*/}.log
exec 2>&1

echo " Customizing Network Parameters "
/usr/sbin/mktcpip -h'nimc1a' -a'9.3.5.13' -m'255.255.254.0' -i'en1'
-g'9.3.4.1' -A'no' -t'N/A'

echo " Customizing Time Zone Settings "
/usr/bin/chtz 'America/Sao_Paulo'

echo " Customizing ulimit for root user"
/usr/bin/chuser fsize=-1 root

exit 0
```

Creating the fb_script resource and listing the resource is shown in Example 3-18.

Example 3-18 Defining and listing a fb_script resource

```
$ nim -o define -t fb_script -a server=master -a
location=/var/adm/ras/fb_script_nimc1ae0.sh fb_script_nimc1ae0
$ lsnim -t fb_script
fb_script_nimc1ae0     resources       fb_script
```

Allocating NIM resources and customizing the client

To perform customization during the installation, we created the resources as shown in Table 3-7 on page 136.

Table 3-7 Resource type and objectives

Resource type	Objective
SPOT and LPP_source	Install the AIX with the desired levels
bosinst_data	Enable TCB, install bundle server and do initiate the installation without prompting for input
fb_script	After creating the desired resources we need to allocate them to the target NIM client machine and start the installation process

Allocating the NIM resources to NIM client nimc1ae0 is shown in Example 3-19:

Example 3-19 Allocating previously resources to the NIM client machine

```
$ nim -o bos_inst -a spot=spot_aix61TL1 -a accept_licenses=yes  -a
lpp_source=lpp_aix61TL1_Full -a bosinst_data=bosinsts_data_nimc1ae0 -a
fb_script=fb_script_nimc1ae0 nimc1ae0
```

The NIM client was assigned and the installation process can be initiated.

The two methods for starting the installation process are:

- Start installation from the NIM master
- Allocate the resources and then boot the NIM client machine, but only in SMS mode (Power System boot mode). This method requires manual data input for the network adapter and for the IP addresses. After initiating the boot the installation begins, as shown in Figure 3-60 on page 137.

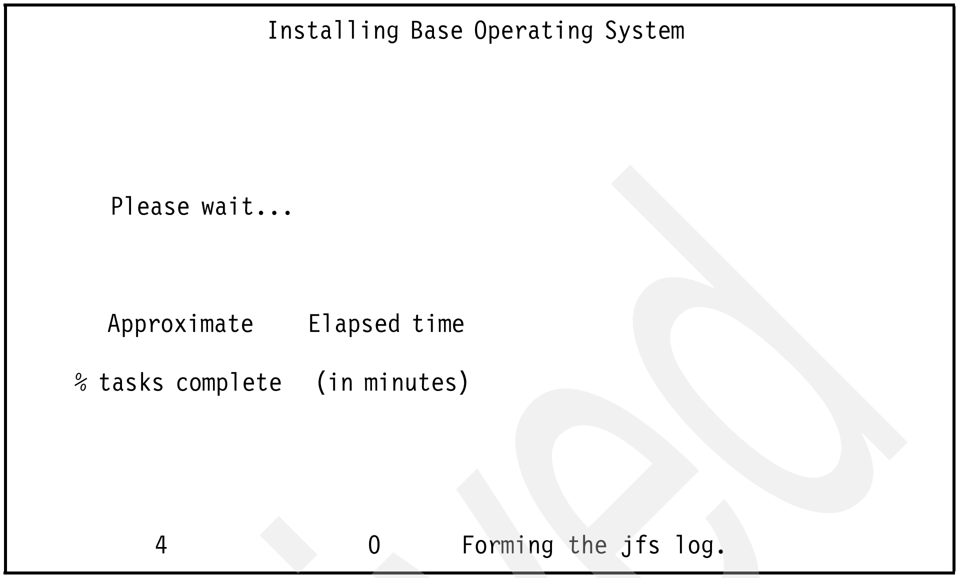

```
                    Installing Base Operating System



        Please wait...


        Approximate    Elapsed time
      % tasks complete  (in minutes)


           4              0      Forming the jfs log.
```

Figure 3-60 NIM installation process

After the system is installed, we can log in to the NIM client and check whether the system was installed with the customized features.

Checking the TCB availability is shown in Example 3-20.

Example 3-20 The tcbck command only works if TCB is enabled

```
# tcbck

Usage: tcbck -a <filename> [ [ <attr> | <attr>=[<value> | ] ] ... ]
             -a -f <filename>
             -a sysck [ treeck_novfs=<dir> ] [ checksum=<program> ]
             -d <filename> | <class> [ <filename> | <class> ] ...
             -d -f <filename>
                 -l /dev/<filename> [ /dev/<filename> ] ...
             -(p|y|n|t) [-io] [ [<filename>|<class>] ... | ALL | tree ]
```

Checking the server bundle is shown in Example 3-21 on page 138.

Example 3-21 Output of lslpp command

```
# lslpp -l |grep -i server
  bos.net.nfs.server        6.1.1.0  COMMITTED  Network File System Server
  bos.net.nis.server        6.1.1.0  COMMITTED  Network Information Service
                                                Server
  bos.net.tcp.server        6.1.1.1  COMMITTED  TCP/IP Server
  devices.vdevice.vty-server.rte
                                                Client/Server Support
  bos.net.nis.server        6.1.1.0  COMMITTED  Network Information Service
                                                Server
  bos.net.tcp.server        6.1.1.1  COMMITTED  TCP/IP Server
  devices.vdevice.vty-server.rte
                                                Client/Server Support
```

Checking network configuration is shown in Example 3-22.

Example 3-22 The netstat and hostname commands showing the customized IP and host name

```
# netstat -rin
Name  Mtu   Network     Address           ZoneID    Ipkts Ierrs    Opkts Oerrs
en0   1500  link#2      a2.4e.57.c6.8.2           1258398     0    19573     0
en0   1500  10.1.1      10.1.1.13                 1258398     0    19573     0
en1   1500  link#3      a2.4e.57.c6.8.3           1264533     0    31262     0
en1   1500  9.3.4       9.3.5.13                  1264533     0    31262     0
lo0   16896 link#1                                    344     0      381     0
lo0   16896 127         127.0.0.1                     344     0      381     0
lo0   16896 ::1                                0      344     0      381     0

# hostname
nimc1a
```

Checking the user limit (**ulimit** command) for root is shown in Example 3-23.

Example 3-23 Output of ulimit -a from the root user

```
# ulimit -a
time(seconds)        unlimited
file(blocks)         unlimited
data(kbytes)         131072
stack(kbytes)        32768
memory(kbytes)       32768
coredump(blocks)     2097151
nofiles(descriptors) 2000
threads(per process) unlimited
processes(per user)  unlimited
```

Customization after installation

In this section, we discuss how to perform post installation customization by using the following NIM resources or options:

- NIM lpp_source resource
- NIM script resources
- Remote or local commands

To perform post installation customization, we can use the NIM *cust* operation.

NIM lpp_source resource

After the installation, this NIM resource can be used to install additional packages that are not required during the installation. It can also be used to install new products or to install fixes. For this example, our NIM client requires the ssh packages that were not installed. See Example 3-24.

Example 3-24 The NIM client nimc1ae0 (host name changed to nimc1a) has no ssh

```
# hostname
nimc1a

# ssh
ksh: ssh:  not found.
```

Define a new lpp_source to provide ssh file sets. Copy the ssh installation files into the `/software/aix-61-exppack/installp/ppc` directory. See Example 3-25.

Example 3-25 The contents of ssh directory after the copy

```
$ ls /software/aix-61-exppack/installp/ppc | grep ssh
openssh.base.4.7.0.5300.I
openssh.license.4.7.0.5300.I
openssh.man.en_US.4.7.0.5300.I
openssh.msg.CA_ES.4.7.0.5300.I
openssh.msg.CS_CZ.4.7.0.5300.I
openssh.msg.EN_US.4.7.0.5300.I
openssh.msg.ES_ES.4.7.0.5300.I
openssh.msg.JA_JP.4.7.0.5300.I
openssh.msg.Ja_JP.4.7.0.5300.I
openssh.msg.KO_KR.4.7.0.5300.I
openssh.msg.PT_BR.4.7.0.5300.I
openssh.msg.ca_ES.4.7.0.5300.I
openssh.msg.cs_CZ.4.7.0.5300.I
openssh.msg.en_US.4.7.0.5300.I
openssh.msg.ja_JP.4.7.0.5300.I
openssh.msg.ko_KR.4.7.0.5300.I
openssh.msg.pt_BR.4.7.0.5300.I
```

Now, we create the new lpp_source for the ssh packages. See Example 3-26.

Example 3-26 Defining a new lpp_source for SSH

```
$ nim -o define -t lpp_source -a server=master -a
location=/software/aix-61-exppack/installp/ppc AIX61_EXPPACK

Preparing to copy install images (this will take several minutes)...

Now checking for missing install images...
warning: 0042-267 c_mk_lpp_source: The defined lpp_source does not have
the "simages" attribute because one or more of the following
packages are missing:
                bos
                bos.net
                bos.diag
                bos.sysmgt
                bos.terminfo
                bos.terminfo.all.data
                devices.graphics
                devices.scsi
                devices.tty
                xlC.rte
                bos.mp64
                devices.common
                bos.64bit
                bos.wpars
                bos.aixpert.cmds
                ifor_ls.base
                ICU4C.rte
                lum.base
                bos.mls
```

To ensure that a new lpp_source can provide bootable resources, the **`c_mk_lpp_source`** command (which is called after issuing the **`nim -o define`**) reads the contents of the file `/usr/lpp/bos.sysmgt/nim/methods/c_sh_lib` to check whether the lpp_source directory contains the required AIX file sets to build a bootable resource.

The file `/usr/lpp/bos.sysmgt/nim/methods/c_sh_lib` is created during the NIM master installation.

> **Important:** Because several file sets are missing, the following message is generated:
>
> ```
> warning: 0042-267 c_mk_lpp_source: The defined lpp_source does not
> have the "simages" attribute because one or more of the following
> packages are missing:
> ```
>
> This message can be safely ignored only when we are creating non-bootable lpp_resources. If we are creating a new lpp_source that will be used during a new installation, any warning message should be investigated before starting the NIM client installation.

Performing the ssh packages installation is shown in Example 3-27.

Example 3-27 Perform a cust NIM operation

```
$ nim -o cust -a lpp_source=AIX61_EXPPACK -a installp_flags="-acgXY"
-a filesets=openssh.base nimc1ae0
```

Verifying the installation log is shown in Example 3-28.

Example 3-28 The nim -o showlog shows the installation log during the NIM client cust

```
$ nim -o showlog -a log_type=niminst nimc1ae0

BEGIN:Tue Sep  9 17:01:01 2008:090920010108
Command line is:
/usr/sbin/installp -acgXY -e /var/adm/ras/nim.installp -f \
/tmp/.workdir.151640.331826_1/.genlib.installp.list.331826-d \
/tmp/_nim_dir_303324/mnt0
...
Successfully updated the Kernel Command Table.
0513-071 The sshd Subsystem has been added.
0513-059 The sshd Subsystem has been started. Subsystem PID is 372942.
Finished processing all filesets.  (Total time:  16 secs).

+-----------------------------------------------------------------------------+
                                Summaries:
+-----------------------------------------------------------------------------+

Installation Summary
--------------------
Name                        Level           Part        Event Result
-------------------------------------------------------------------------------
```

```
openssl.base                0.9.8.600       USR         APPLY SUCCESS
openssl.base                0.9.8.600       ROOT        APPLY SUCCESS
openssh.base.client         4.7.0.5300      USR         APPLY SUCCESS
openssh.base.server         4.7.0.5300      USR         APPLY SUCCESS
openssh.base.client         4.7.0.5300      ROOT        APPLY SUCCESS
openssh.base.server         4.7.0.5300      ROOT        APPLY SUCCESS

END:Tue Sep  9 17:01:17 2008:090920011708
```

Log in on the NIM client to verify the ssh installation, as shown in Example 3-29.

Example 3-29 NIM client nimc1a after the customization operation

```
# hostname
nimc1a

# ssh
usage: ssh [-1246AaCfgKkMNnqsTtVvXxY] [-b bind_address] [-c cipher_spec]
           [-D [bind_address:]port] [-e escape_char] [-F configfile]
           [-i identity_file] [-L [bind_address:]port:host:hostport]
           [-l login_name] [-m mac_spec] [-O ctl_cmd] [-o option] [-p port]
           [-R [bind_address:]port:host:hostport] [-S ctl_path]
           [-w local_tun[:remote_tun]] [user@]hostname [command]
```

NIM script resources

This resource can perform any additional customization by using a regular shell script. We used this resource to create new users and groups, add new `/etc/hosts`, and configure the `/etc/motd` file.

Creating a script is shown in Example 3-30.

Example 3-30 Shows a simple script created to customize the nimc1ae0

```
# !/bin/ksh
# Script customize_script_nimc1ae0.sh
# This script was created to customize the NIM Client nimc1ae0
#
# Created in 2008
# Objective : Customize TimeZone, ulimit and Network Information
#
# Redirect the STDOUT and STDERR to a logfile
exec 1> /tmp/${0##*/}.log
exec 2>&1

echo " #### Adding extra entries do /etc/hosts #### "
cp /etc/hosts /etc/hosts.`date "+%h_%d_%y"`
```

```
cat <<EOF >> /etc/hosts
9.3.5.3         vio1a
9.3.5.4         vio2a
9.3.5.5         nim1a
9.3.5.6         uam1a
9.3.5.7         tpm1a
9.3.5.8         tbo1a
9.3.5.9         wpar1a
9.3.5.10        itm1a
9.3.5.11        tsm1a
9.3.5.13        nimd1a
9.3.5.128       hmc1a

EOF

echo " #### Adding Users and Groups  #### "
mkgroup tengrp
sleep 5
for i in 1 2 3 4 5 6 7 8 9
do
mkuser pgrp=tengrp tenant$i
done

echo " #### Customizing /etc/motd file #### "

cp /etc/motd /etc/motd.`date "+%h_%d_%y"`
cat <<EOF > /etc/motd

This Server is under Company A administration.
The access to this Server is restricted to authorized
users only.

EOF

exit 0
```

Defining a resource type *script* is shown in Example 3-31.

Example 3-31   Defining a script resource using command line

```
$ nim -o define -t script -a server=master -a
location=/var/adm/ras/customize_script_nimc1ae0.sh cust_script_nimc1ae0
```

Listing the created resource is shown in Example 3-32.

Example 3-32 lsnim command

```
$ lsnim -t script
cust_script_nimc1ae0     resources       script
```

Allocating NIM script and customizing the client

The script resources in Table 3-8 have to be allocated to the NIM client before we can perform the customization operation.

Table 3-8 The definition of script NIM resource

Type of resource	Description
script	Customize the /etc/motd file, create users and group and add entries in /etc/hosts file

Allocating the NIM resources to NIM client nimc1ae0 is shown in Example 3-33.

Example 3-33 nim -o command to allocate and customize the client

```
$ nim -o allocate -a script=cust_script_nimc1ae0 nimc1ae0

$ nim -o cust nimc1ae0
```

After the customize operation (cust) is finished, we can log in to the NIM client nimc1ae0 to verify the customized attributes.

Log in on nimclient to check customized /etc/motd file. The result is shown in Figure 3-61.

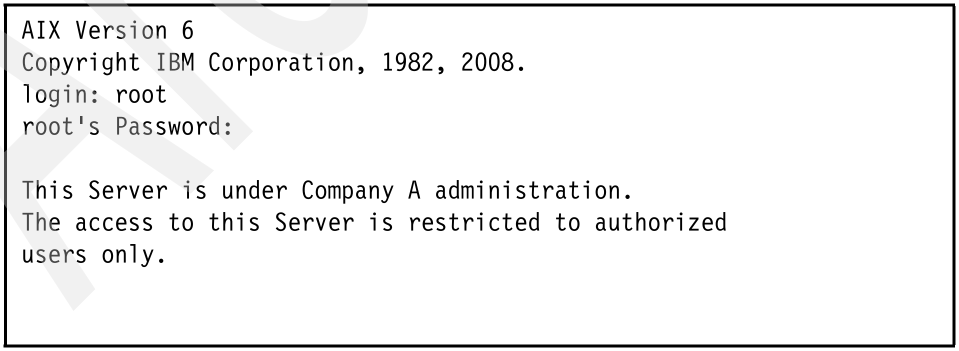

```
AIX Version 6
Copyright IBM Corporation, 1982, 2008.
login: root
root's Password:

This Server is under Company A administration.
The access to this Server is restricted to authorized
users only.
```

Figure 3-61 Panel shows the motd customized warning

Verifying the new users is shown in Example 3-34.

Example 3-34 /etc/passwd file after customization

```
# cat /etc/passwd
root:!:0:0::/:/usr/bin/ksh
daemon:!:1:1::/etc:
bin:!:2:2::/bin:
sys:!:3:3::/usr/sys:
adm:!:4:4::/var/adm:
uucp:!:5:5::/usr/lib/uucp:
guest:!:100:100::/home/guest:
nobody:!:4294967294:4294967294::/:
lpd:!:9:4294967294::/:
lp:*:11:11::/var/spool/lp:/bin/false
invscout:*:6:12::/var/adm/invscout:/usr/bin/ksh
snapp:*:200:13:snapp login user:/usr/sbin/snapp:/usr/sbin/snappd
ipsec:*:201:1::/etc/ipsec:/usr/bin/ksh
nuucp:*:7:5:uucp login user:/var/spool/uucppublic:/usr/sbin/uucp/uucico
pconsole:*:8:0::/var/adm/pconsole:/usr/bin/ksh
esaadmin:*:10:0::/var/esa:/usr/bin/ksh
sshd:*:202:201::/var/empty:/usr/bin/ksh
tenant1:*:203:202::/home/tenant1:/usr/bin/ksh
tenant2:*:204:202::/home/tenant2:/usr/bin/ksh
tenant3:*:205:202::/home/tenant3:/usr/bin/ksh
tenant4:*:206:202::/home/tenant4:/usr/bin/ksh
tenant5:*:207:202::/home/tenant5:/usr/bin/ksh
tenant6:*:208:202::/home/tenant6:/usr/bin/ksh
tenant7:*:209:202::/home/tenant7:/usr/bin/ksh
tenant8:*:210:202::/home/tenant8:/usr/bin/ksh
tenant9:*:211:202::/home/tenant9:/usr/bin/ksh
```

Verifying the new group is shown in Example 3-35.

Example 3-35 Verifying /etc/group customized file

```
# cat /etc/group
system:!:0:root,pconsole,esaadmin
staff:!:1:ipsec,esaadmin,sshd
bin:!:2:root,bin
sys:!:3:root,bin,sys
adm:!:4:bin,adm
uucp:!:5:uucp,nuucp
mail:!:6:
security:!:7:root
cron:!:8:root
```

```
printq:!:9:lp
audit:!:10:root
ecs:!:28:
nobody:!:4294967294:nobody,lpd
usr:!:100:guest
perf:!:20:
shutdown:!:21:
lp:!:11:root,lp
invscout:!:12:invscout
snapp:!:13:snapp
ipsec:!:200:
pconsole:!:14:pconsole
sshd:!:201:sshd
tengrp:!:202:tenant1,tenant2,tenant3,tenant4,tenant5,tenant6,tenant7,te
nant8,tenant9
```

Verifying customized `/etc/hosts` file is shown in Example 3-36.

Example 3-36 Contents of /etc/hosts file

```
# Internet Address      Hostname        # Comments
# 192.9.200.1           net0sample      # ethernet name/address
# 128.100.0.1           token0sample    # token ring name/address
# 10.2.0.2              x25sample       # x.25 name/address
# 2000:1:1:1:209:6bff:feee:2b7f ipv6sample      # ipv6 name/address
127.0.0.1 loopback localhost      # loopback (lo0) name/address
10.1.1.5        nim1ae0
9.3.5.13        nimc1a
9.3.5.3         vio1a
9.3.5.4         vio2a
9.3.5.5         nim1a
9.3.5.6         uam1a
9.3.5.7         tpm1a
9.3.5.8         tbo1a
9.3.5.9         wpar1a
9.3.5.10        itm1a
9.3.5.11        tsm1a
9.3.5.13        nimd1a
9.3.5.128       hmc1a
```

Manual and remote commands

Any customization can be done locally at the NIM client by the command line, the **`smit`** command, or the **`remote`** command.

We can use the rexec remote commands to start a command from the NIM master to the NIM client.

> **Note:** To perform the **`rexec`** command, the NIM client machine has to be properly configured. For further information refer to the **`netrc`** man page.

3.7.2 Customization with Tivoli Provisioning Manager

This section discusses customization options that can be performed by using TPM. TPM provides most customization features with a GUI.

Customizing network settings

Use TPM to customize network parameters such as host name and IP address (and gateway and netmask).

The host name

TPM can change the host name when we create the virtual server. When we create the virtual server, we can allocate host name in Name field, as shown in Figure 3-62.

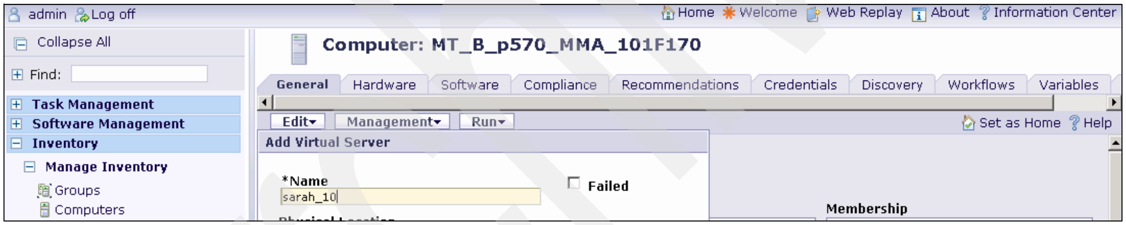
Computer: MT_B_p570_MMA_101F170

Add Virtual Server

*Name: sarah_10

☐ Failed

Figure 3-62 Allocate host name in virtual partition when add virtual server

IP address, gateway, netmask

TPM can change IP address, gateway, netmask after we created virtual server.

After creating the virtual server, we select **Inventory** → **Manage Inventory** → **Computers**, and then select the computer we created. Then, select the **Hardware** tab.

To view the network resources, select **Add Interface**. See Figure 3-63 on page 148.

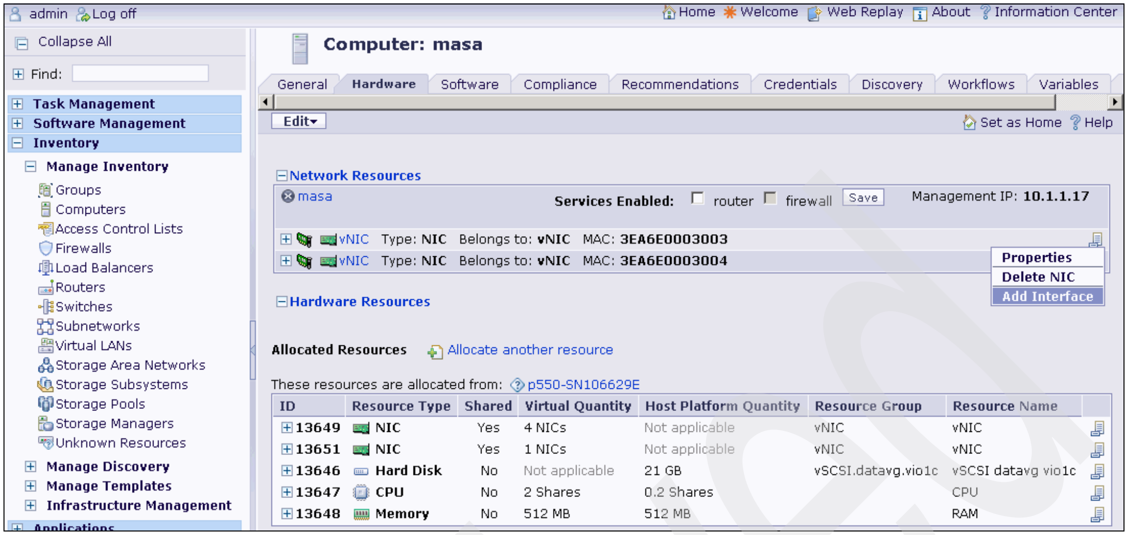

Network Resources

masa Services Enabled: □ router □ firewall Save Management IP: 10.1.1.17

vNIC Type: NIC Belongs to: vNIC MAC: 3EA6E0003003

vNIC Type: NIC Belongs to: vNIC MAC: 3EA6E0003004

Hardware Resources

Allocated Resources

These resources are allocated from: p550-SN106629E

ID	Resource Type	Shared	Virtual Quantity	Host Platform Quantity	Resource Group	Resource Name
13649	NIC	Yes	4 NICs	Not applicable	vNIC	vNIC
13651	NIC	Yes	1 NICs	Not applicable	vNIC	vNIC
13646	Hard Disk	No	Not applicable	21 GB	vSCSI.datavg.vio1c	vSCSI datavg vio1c
13647	CPU	No	2 Shares	0.2 Shares		CPU
13648	Memory	No	512 MB	512 MB		RAM

Figure 3-63 Add interface in partition's hardware tab

In **Add Interface**, we can add a network interface, as shown in Figure 3-64 on page 149.

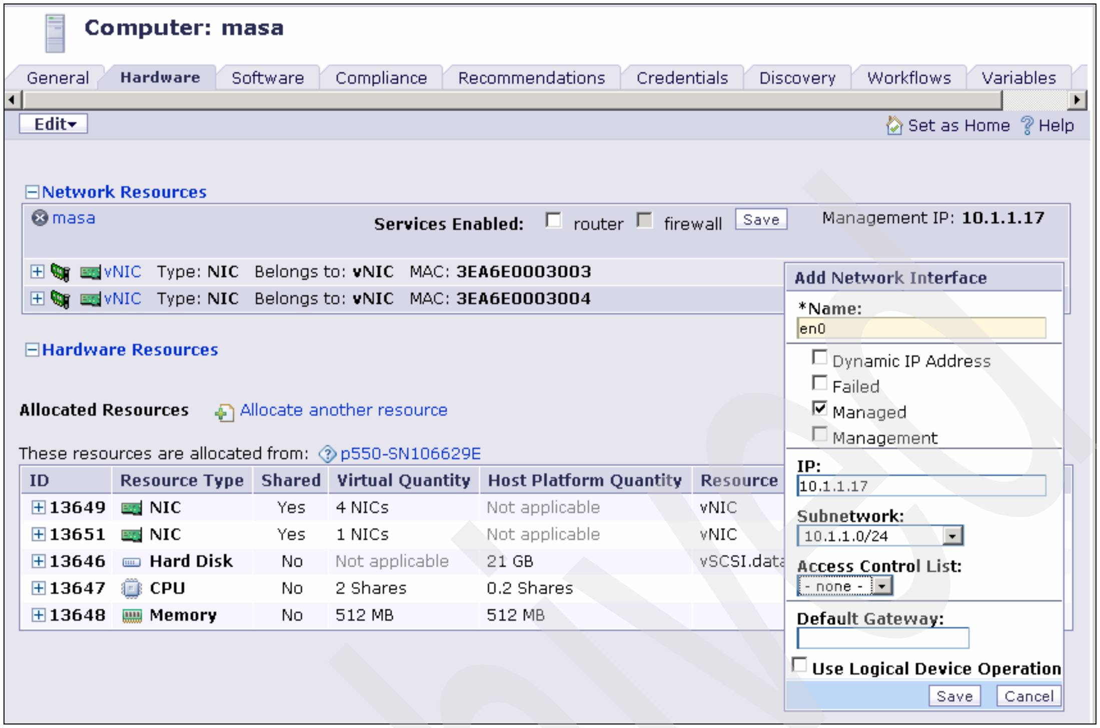

masa Services Enabled: □ router □ firewall Save Management IP: **10.1.1.17**

vNIC Type: **NIC** Belongs to: **vNIC** MAC: **3EA6E0003003**

vNIC Type: **NIC** Belongs to: **vNIC** MAC: **3EA6E0003004**

Allocated Resources

These resources are allocated from: p550-SN106629E

ID	Resource Type	Shared	Virtual Quantity	Host Platform Quantity	Resource
13649	NIC	Yes	4 NICs	Not applicable	vNIC
13651	NIC	Yes	1 NICs	Not applicable	vNIC
13646	Hard Disk	No	Not applicable	21 GB	vSCSI.data
13647	CPU	No	2 Shares	0.2 Shares	
13648	Memory	No	512 MB	512 MB	

Add Network Interface

*Name: en0

□ Dynamic IP Address
□ Failed
☑ Managed
□ Management

IP: 10.1.1.17

Subnetwork: 10.1.1.0/24

Access Control List: - none -

Default Gateway:

□ Use Logical Device Operation

Figure 3-64 Add network interface in partition's hardware tab

After adding the interface, expand the network device name in **Network Resources**. Then, change the network information in the **Properties** menu, as shown in Figure 3-65 on page 150.

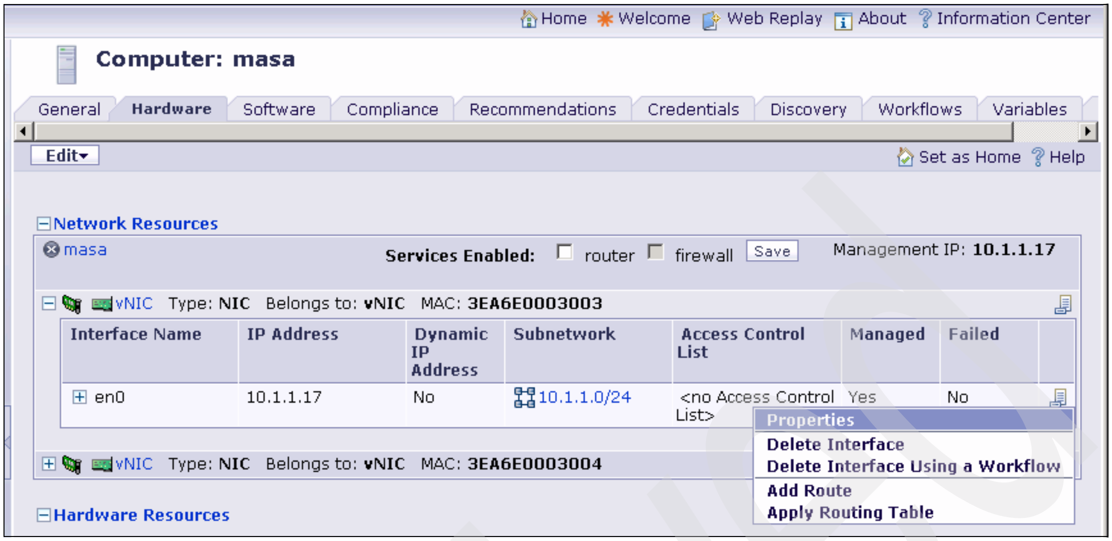

Computer: masa

Network Resources

masa Services Enabled: □ router □ firewall Management IP: 10.1.1.17

vNIC Type: NIC Belongs to: vNIC MAC: 3EA6E0003003

Interface Name	IP Address	Dynamic IP Address	Subnetwork	Access Control List	Managed	Failed
en0	10.1.1.17	No	10.1.1.0/24	<no Access Control List>	Yes	No

vNIC Type: NIC Belongs to: vNIC MAC: 3EA6E0003004

Hardware Resources

Figure 3-65 Properties menu to change network information

Select **Properties** to change the IP and subnetwork mask, as shown in Figure 3-66 on page 151.

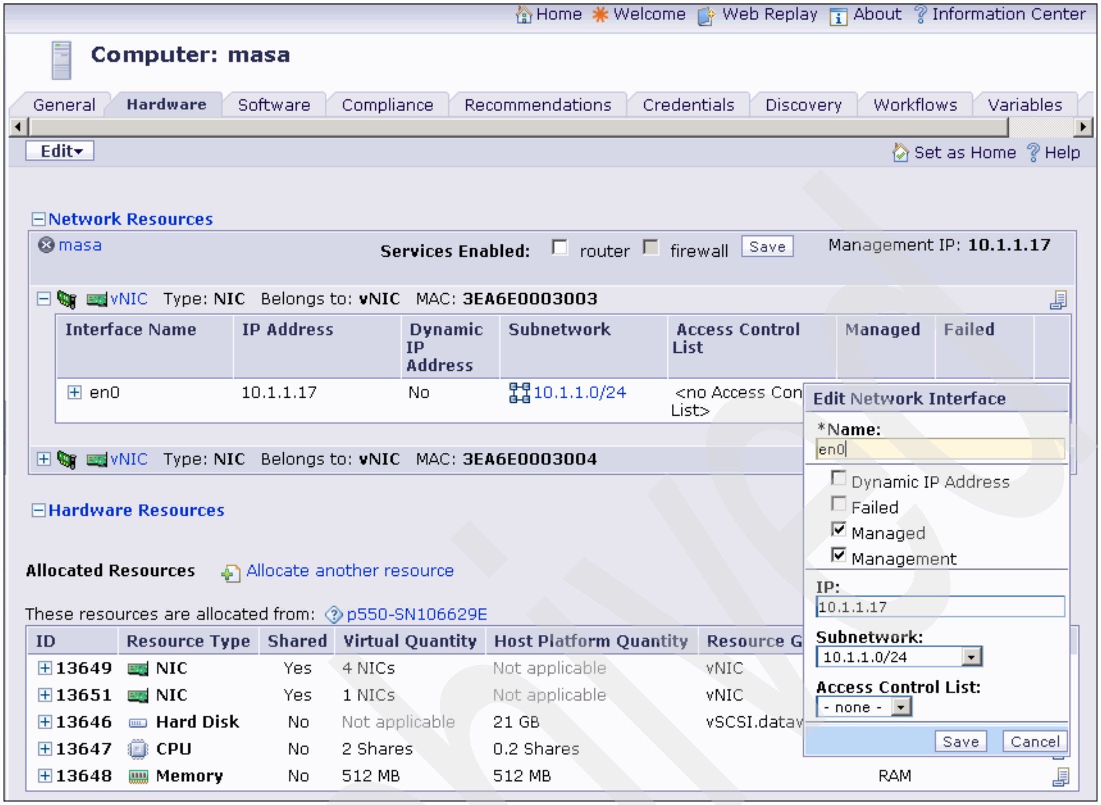

Computer: masa

Network Resources

masa Services Enabled: □ router □ firewall Save Management IP: **10.1.1.17**

vNIC Type: **NIC** Belongs to: **vNIC** MAC: **3EA6E0003003**

Interface Name	IP Address	Dynamic IP Address	Subnetwork	Access Control List	Managed	Failed
en0	10.1.1.17	No	10.1.1.0/24	<no Access Con List>		

vNIC Type: **NIC** Belongs to: **vNIC** MAC: **3EA6E0003004**

Hardware Resources

Allocated Resources Allocate another resource

These resources are allocated from: p550-SN106629E

ID	Resource Type	Shared	Virtual Quantity	Host Platform Quantity	Resource G	
13649	NIC	Yes	4 NICs	Not applicable	vNIC	
13651	NIC	Yes	1 NICs	Not applicable	vNIC	
13646	Hard Disk	No	Not applicable	21 GB	vSCSI.datav	
13647	CPU	No	2 Shares	0.2 Shares		
13648	Memory	No	512 MB	512 MB		RAM

Edit Network Interface

*Name: en0

□ Dynamic IP Address
□ Failed
☑ Managed
☑ Management

IP: 10.1.1.17

Subnetwork: 10.1.1.0/24

Access Control List: - none -

Save Cancel

Figure 3-66 Change IP and Subnetwork mask in network interface

Customizing users and groups

To create user or group With TPM, we can send AIX command directly to a specific computer.

First, we have to determine the device ID of virtual partition. See Figure 3-19 on page 86 and Figure 3-20 on page 87. After you determine the device ID of virtual partition, send the command directly to virtual server from TPM.

Users

In the Find field, enter `Device.ExecuteCommand` and then click `Device.ExecuteCommand`. The Workflow Editor is displayed. See Figure 3-67 on page 152.

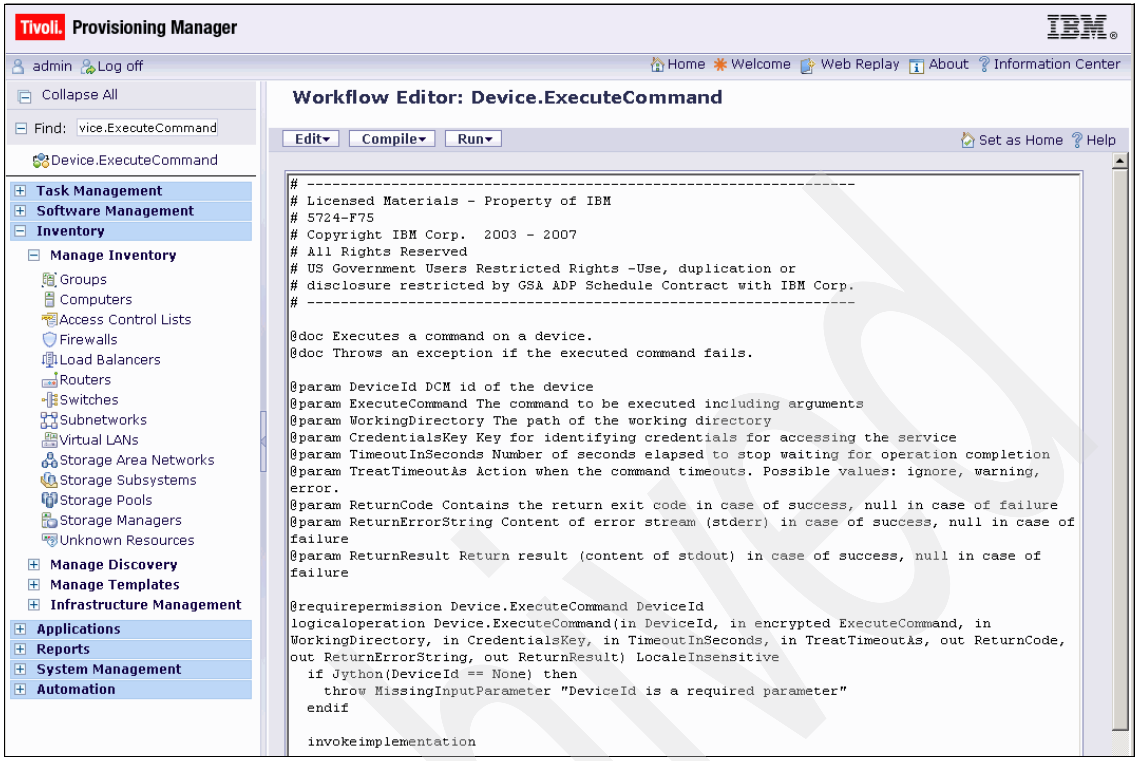

```
# ---------------------------------------------------------------------
# Licensed Materials - Property of IBM
# 5724-F75
# Copyright IBM Corp.  2003 - 2007
# All Rights Reserved
# US Government Users Restricted Rights -Use, duplication or
# disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
# ---------------------------------------------------------------------

@doc Executes a command on a device.
@doc Throws an exception if the executed command fails.

@param DeviceId DCM id of the device
@param ExecuteCommand The command to be executed including arguments
@param WorkingDirectory The path of the working directory
@param CredentialsKey Key for identifying credentials for accessing the service
@param TimeoutInSeconds Number of seconds elapsed to stop waiting for operation completion
@param TreatTimeoutAs Action when the command timeouts. Possible values: ignore, warning,
error.
@param ReturnCode Contains the return exit code in case of success, null in case of failure
@param ReturnErrorString Content of error stream (stderr) in case of success, null in case of
failure
@param ReturnResult Return result (content of stdout) in case of success, null in case of
failure

@requirepermission Device.ExecuteCommand DeviceId
logicaloperation Device.ExecuteCommand(in DeviceId, in encrypted ExecuteCommand, in
WorkingDirectory, in CredentialsKey, in TimeoutInSeconds, in TreatTimeoutAs, out ReturnCode,
out ReturnErrorString, out ReturnResult) LocaleInsensitive
  if Jython(DeviceId == None) then
    throw MissingInputParameter "DeviceId is a required parameter"
  endif

  invokeimplementation
```

Figure 3-67 ExecuteCommand workflow editor

Click **Run**. In the DeviceID field, enter the device ID of the virtual partition. In the ExecuteCommand field, enter `mkuser sarah` as shown in Figure 3-68. Then, click **Run**. We can create the user named `sarah` with the default values.

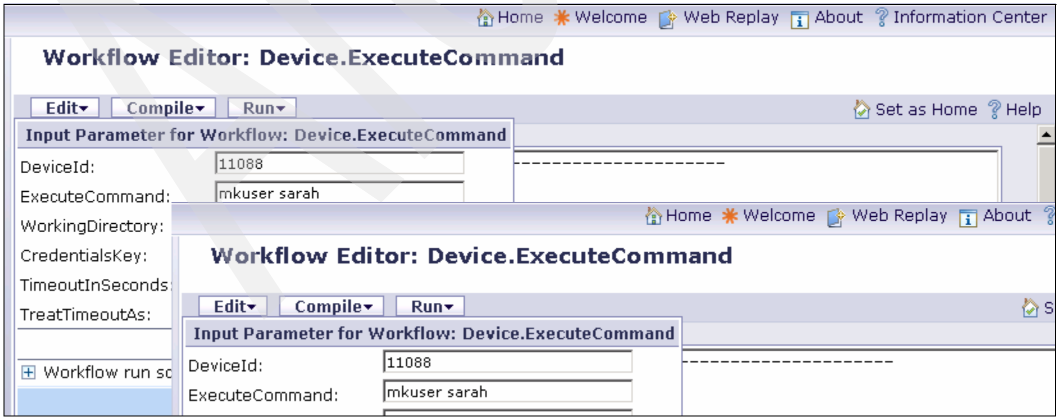

DeviceId: 11088

ExecuteCommand: mkuser sarah

Figure 3-68 Create user with Device.ExecuteCommand directly

The result are displayed in the virtual server, as shown in Example 3-37.

Example 3-37 User created in virtual server

```
# cat /etc/passwd
root:!:0:0::/root:/usr/bin/ksh
daemon:!:1:1::/etc:
bin:!:2:2::/bin:
sys:!:3:3::/usr/sys:
adm:!:4:4::/var/adm:
uucp:!:5:5::/usr/lib/uucp:
guest:!:100:100::/home/guest:
nobody:!:4294967294:4294967294::/:
lpd:!:9:4294967294::/:
lp:*:11:11::/var/spool/lp:/bin/false
invscout:*:6:12::/var/adm/invscout:/usr/bin/ksh
snapp:*:200:13:snapp login user:/usr/sbin/snapp:/usr/sbin/snappd
nuucp:*:7:5:uucp login user:/var/spool/uucppublic:/usr/sbin/uucp/uucico
ipsec:*:201:1::/etc/ipsec:/usr/bin/ksh
sshd:*:202:201::/var/empty:/usr/bin/ksh
tuam:*:203:202::/home/tuam:/usr/bin/ksh
sarah:*:206:1::/home/sarah:/usr/bin/ksh
#
```

Group

To create a group, use same method as we used to create a user in the virtual server.

Click **Run**. In the DeviceID field, enter the device ID of the virtual partition. In the ExecuteCommand field, enter `mkgroup finance` as shown in Figure 3-69 on page 154. Then, click **Run**. We can create the group `finance` with the default values.

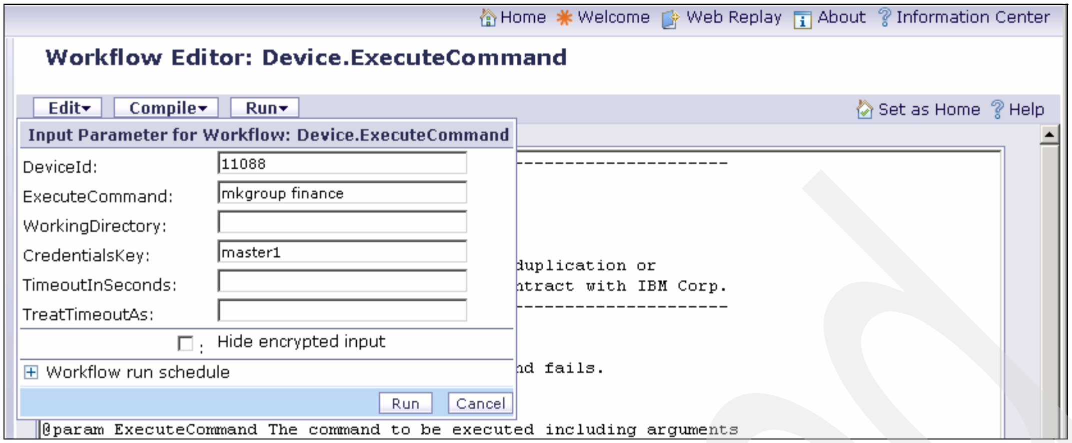

Workflow Editor: Device.ExecuteCommand

Input Parameter for Workflow: Device.ExecuteCommand

DeviceId: 11088
ExecuteCommand: mkgroup finance
WorkingDirectory:
CredentialsKey: master1
TimeoutInSeconds:
TreatTimeoutAs:

□ : Hide encrypted input

⊞ Workflow run schedule

```
---------------------
duplication or
ntract with IBM Corp.
---------------------
nd fails.
@param ExecuteCommand The command to be executed including arguments
```

Figure 3-69 Create group with Device.ExecuteCommand directly

The results in the virtual server are shown in Example 3-38.

Example 3-38 Group created in virtual server

```
# cat /etc/group
system:!:0:root
staff:!:1:ipsec,sshd,tuam,sarah
bin:!:2:root,bin
sys:!:3:root,bin,sys
adm:!:4:bin,adm
uucp:!:5:uucp,nuucp
mail:!:6:
security:!:7:root
cron:!:8:root
printq:!:9:lp
audit:!:10:root
ecs:!:28:
nobody:!:4294967294:nobody,lpd
usr:!:100:guest
perf:!:20:
shutdown:!:21:
lp:!:11:root,lp
invscout:!:12:invscout
snapp:!:13:snapp
ipsec:!:200:
sshd:!:201:sshd
tuam:!:202:tuam
finance:!:205:
#
```

Customizing the environment configuration

TPM provides environment configuration during the installing of the operating system or software.

The TimeZone setting

Before OS installation, we might want to set the time zone for the target server. Select **Software Management** → **Manage Software Catalog** → **Images**. Select boot server, and then select **Edit** → **Install Image**.

Select a value for the TimeZone parameter. See Figure 3-70.

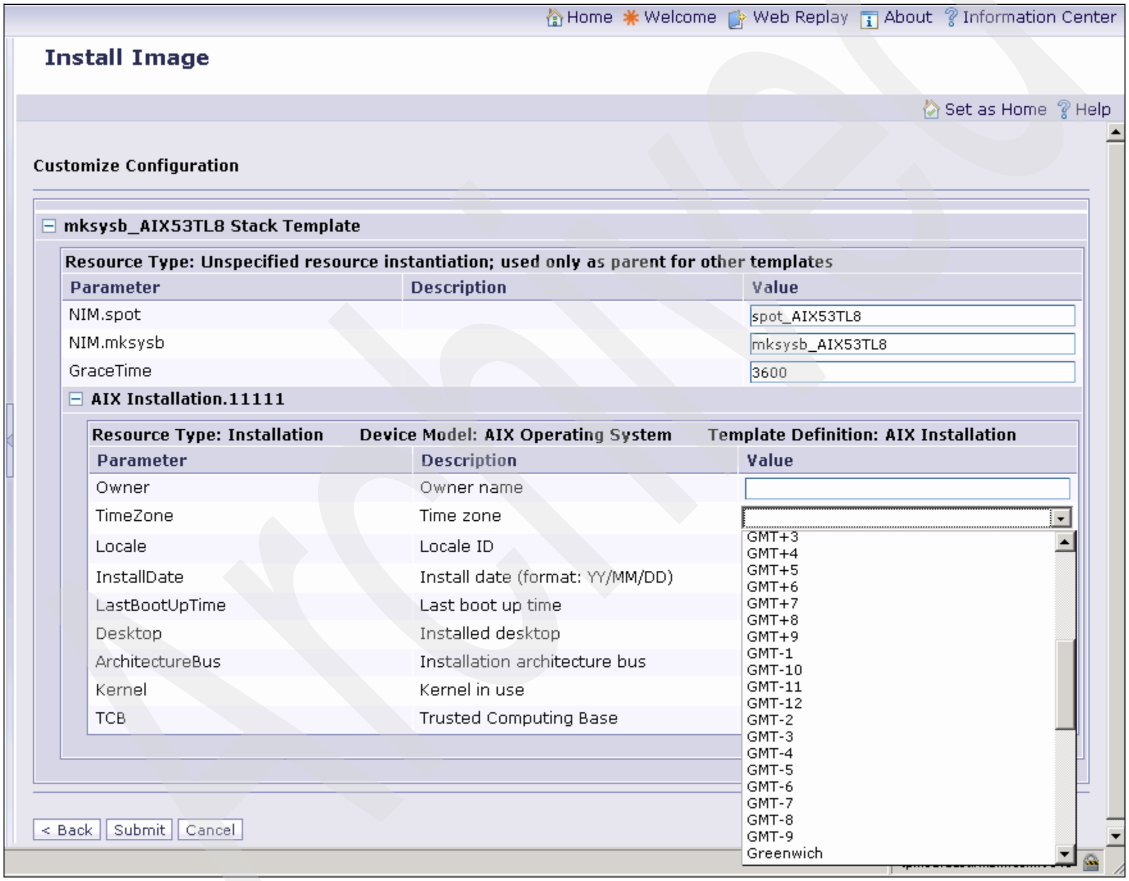

Install Image

Customize Configuration

⊟ **mksysb_AIX53TL8 Stack Template**

Resource Type: Unspecified resource instantiation; used only as parent for other templates

Parameter	Description	Value
NIM.spot		spot_AIX53TL8
NIM.mksysb		mksysb_AIX53TL8
GraceTime		3600

⊟ **AIX Installation.11111**

Resource Type: Installation **Device Model: AIX Operating System** **Template Definition: AIX Installation**

Parameter	Description	Value
Owner	Owner name	
TimeZone	Time zone	
Locale	Locale ID	
InstallDate	Install date (format: YY/MM/DD)	
LastBootUpTime	Last boot up time	
Desktop	Installed desktop	
ArchitectureBus	Installation architecture bus	
Kernel	Kernel in use	
TCB	Trusted Computing Base	

GMT+3
GMT+4
GMT+5
GMT+6
GMT+7
GMT+8
GMT+9
GMT-1
GMT-10
GMT-11
GMT-12
GMT-2
GMT-3
GMT-4
GMT-5
GMT-6
GMT-7
GMT-8
GMT-9
Greenwich

Figure 3-70 TimeZone customization in the Install Image window

Locale

Before OS installation, you might want to set a locale for the target server. Select **Software Management** → **Manage Software Catalog** → **Images**. Select boot server, and then select **Edit** → **Install Image**.

Select a value for the Locale parameter, as shown in Figure 3-71.

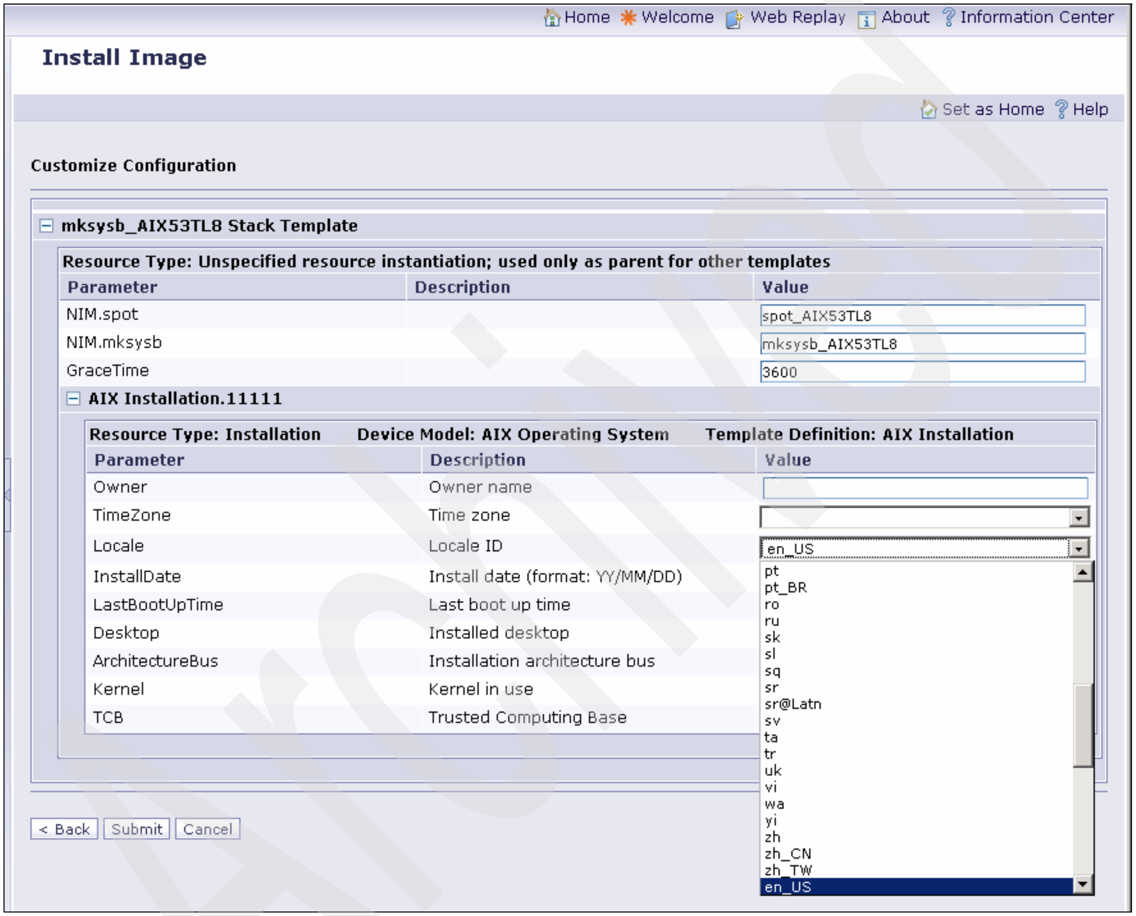

Install Image

Customize Configuration

mksysb_AIX53TL8 Stack Template

Resource Type: Unspecified resource instantiation; used only as parent for other templates

Parameter	Description	Value
NIM.spot		spot_AIX53TL8
NIM.mksysb		mksysb_AIX53TL8
GraceTime		3600

AIX Installation.11111

Resource Type: Installation Device Model: AIX Operating System Template Definition: AIX Installation

Parameter	Description	Value
Owner	Owner name	
TimeZone	Time zone	
Locale	Locale ID	en_US
InstallDate	Install date (format: YY/MM/DD)	
LastBootUpTime	Last boot up time	
Desktop	Installed desktop	
ArchitectureBus	Installation architecture bus	
Kernel	Kernel in use	
TCB	Trusted Computing Base	

Figure 3-71 Locale customization in Install Image window

After selecting all customization options, click **Submit**. The OS is installed as configured by using the settings from the dialog previously mentioned.

Customizing the user environment and files

TPM provides simple software product distribution to upload customized file.

To customize `/etc/environment`, `/etc/motd` and `/etc/profile`, we make our own customized file, adding our own environment variable. Then, we create a tar file with the file we edited, copy the tar file to `%TIO_HOME/repository/package` directory, as shown in Example 3-39.

> **Note:** We can use any name instead of `package`, but make sure to create it under `%TIO_HOME/repository` directory.

Example 3-39 Packaging customized file in one tar file

```
$ mkdir package
$ cd package
$ pwd
/opt/ibm/tivoli/tpm/repository/package
.... editing file ...
$ tar cvf custom.tar motd environment profile

$ ls -la
total 32
drwxr-xr-x    2 tioadmin tioadmin        256 Sep 12 16:16 .
drwxrwxr-x   53 tioadmin tioadmin       4096 Sep 12 16:16 ..
-rw-r--r--    1 tioadmin tioadmin      10240 Sep 12 16:15 custom.tar
$
$ tar tvf custom.tar
-r--r--r-- bin/bin         880 2008-09-08 00:34:02 motd
-rw-rw-r-- root/system    1915 2008-09-08 08:26:55 environment
-r-xr-xr-x bin/bin        1801 2008-09-08 00:34:03 profile
$
```

Select **Software Management** → **Manage Software Catalog** → **Software Products**. Then, select **Edit** → **Add Software Product** to add a new software definition to the software catalog, as shown in Figure 3-72 on page 158.

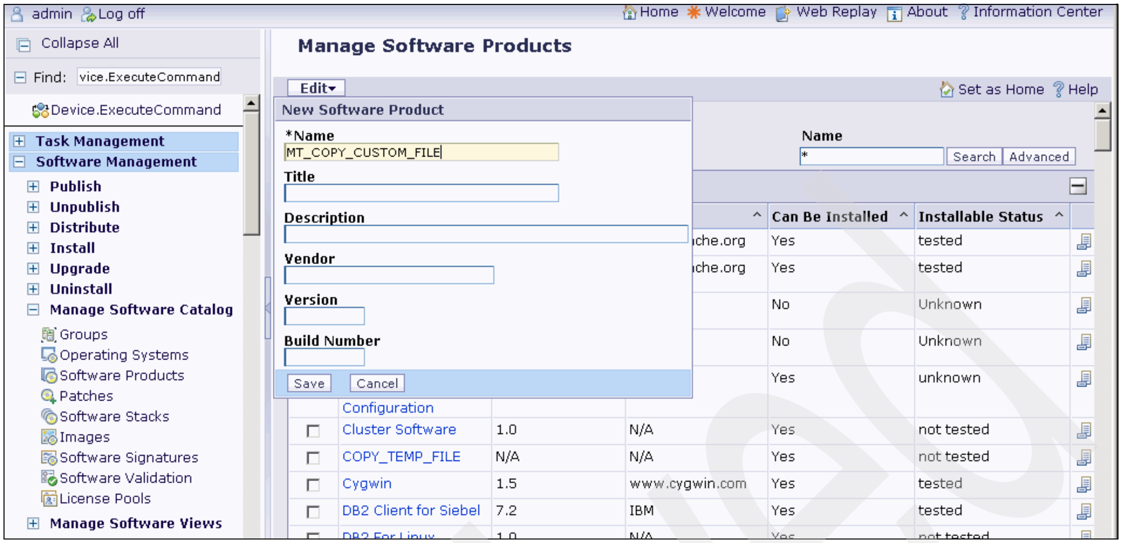

New Software Product

*Name: MT_COPY_CUSTOM_FILE

Title:

Description:

Vendor:

Version:

Build Number:

	Name	Version	Vendor	Can Be Installed	Installable Status
			ache.org	Yes	tested
			ache.org	Yes	tested
				No	Unknown
				No	Unknown
				Yes	unknown
	Configuration				
□	Cluster Software	1.0	N/A	Yes	not tested
□	COPY_TEMP_FILE	N/A	N/A	Yes	not tested
□	Cygwin	1.5	www.cygwin.com	Yes	tested
□	DB2 Client for Siebel	7.2	IBM	Yes	tested
□	DB2 For Linux	1.0	N/A	Yes	not tested

Figure 3-72 Creating software product in Manage Software Product menu

We select the name of the software product we created by selecting **Edit** → **Add Installable**, as shown in Figure 3-73.

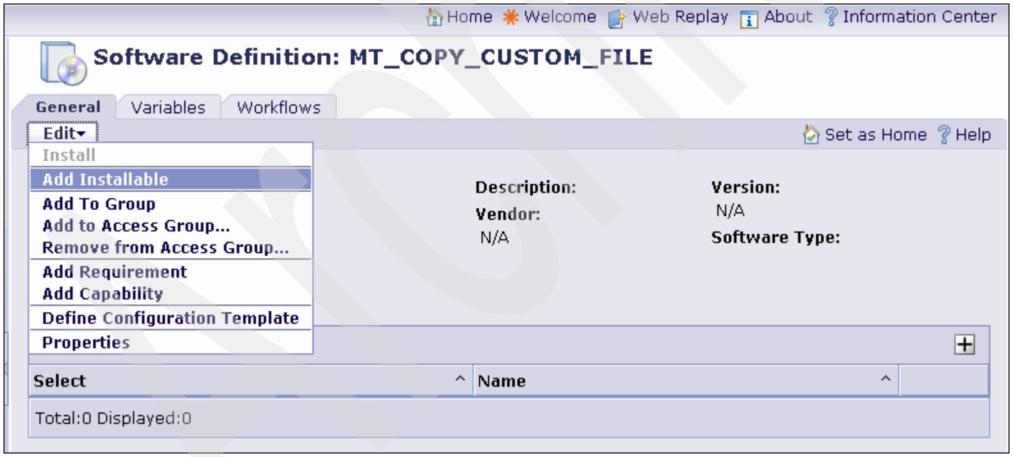

Software Definition: MT_COPY_CUSTOM_FILE

Description:

Vendor: N/A

Version: N/A

Software Type:

Select	Name
Total:0 Displayed:0	

Figure 3-73 Creating Software Installable in Software definition

Specify the required information. For the File Repository, select **LocalFileRepository**. For Installable Path, enter the directory name we created under `%TIO_HOME/repository`. For File Name, enter the tar file name as shown in Figure 3-74 on page 159.

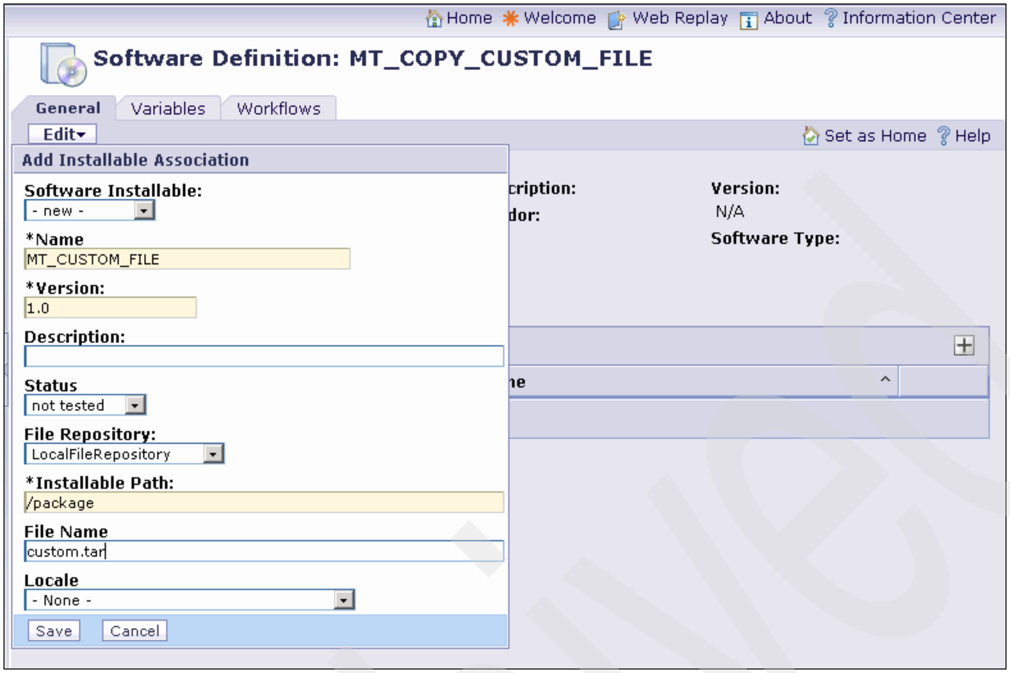

Software Definition: MT_COPY_CUSTOM_FILE

Add Installable Association

Software Installable: - new -

*Name: MT_CUSTOM_FILE

*Version: 1.0

Description:

Status: not tested

File Repository: LocalFileRepository

*Installable Path: /package

File Name: custom.tar

Locale: - None -

cription:

dor:

Version: N/A

Software Type:

Figure 3-74 Setting the parameter for installing the packaged file

Create a software configuration template for the software product. Select **Edit** → **Define Configuration Template** as shown in Figure 3-75 on page 160.

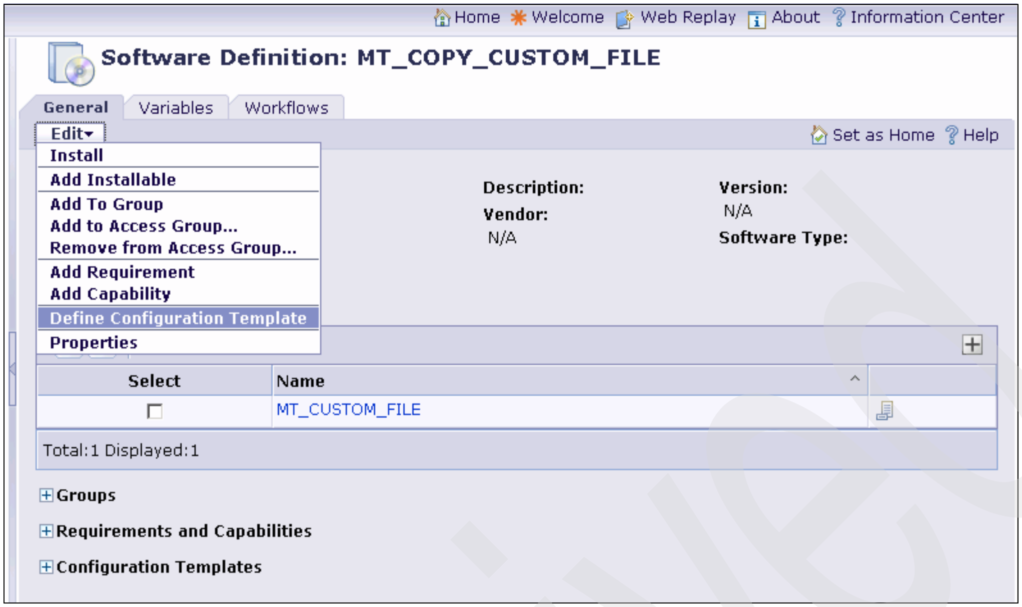

Description:

Vendor: N/A

Version: N/A

Software Type:

Select	Name
□	MT_CUSTOM_FILE

Total:1 Displayed:1

Figure 3-75   Create a software configuration template for the software product

For Creation Method, select **Create from a template definition**. For Configuration Template Definition, select **Hosting-Environment:AIX: Extract Only** as shown in Figure 3-76.

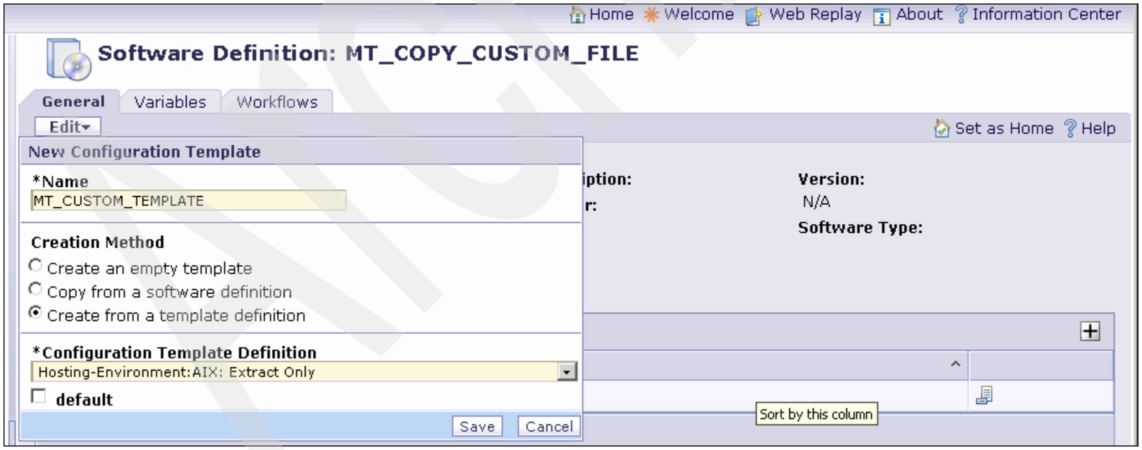

New Configuration Template

*Name: MT_CUSTOM_TEMPLATE

Creation Method
○ Create an empty template
○ Copy from a software definition
● Create from a template definition

*Configuration Template Definition: Hosting-Environment:AIX: Extract Only

□ default

iption:

r:

Version: N/A

Software Type:

Figure 3-76   Setting a parameter for configuration template definition

After the template is created, expand the **Configuration Templates** and the template name we created. Beside the template name, right-click the icon and select **Edit** as shown in Figure 3-77 on page 161.

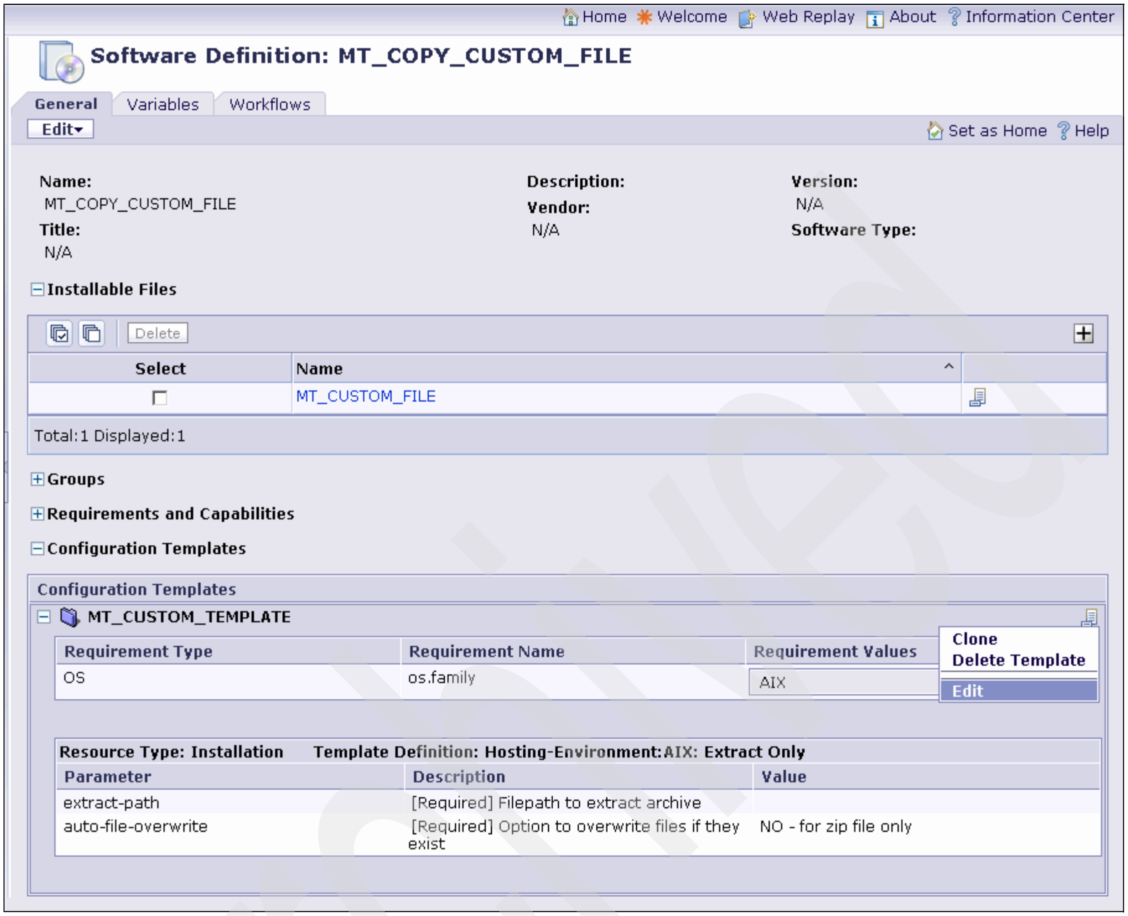

Software Definition: MT_COPY_CUSTOM_FILE

Name: MT_COPY_CUSTOM_FILE
Title: N/A
Description:
Vendor: N/A
Version: N/A
Software Type:

⊟ Installable Files

Select	Name
□	MT_CUSTOM_FILE

Total:1 Displayed:1

⊞ Groups
⊞ Requirements and Capabilities
⊟ Configuration Templates

Configuration Templates
⊟ MT_CUSTOM_TEMPLATE

Requirement Type	Requirement Name	Requirement Values
OS	os.family	AIX

Resource Type: Installation   Template Definition: Hosting-Environment:AIX: Extract Only

Parameter	Description	Value
extract-path	[Required] Filepath to extract archive	
auto-file-overwrite	[Required] Option to overwrite files if they exist	NO - for zip file only

Figure 3-77 Editing a template parameter for installation

For the extract-path, enter `/etc` and for auto-file-overwrite, select **YES** as shown in Figure 3-78 on page 162.

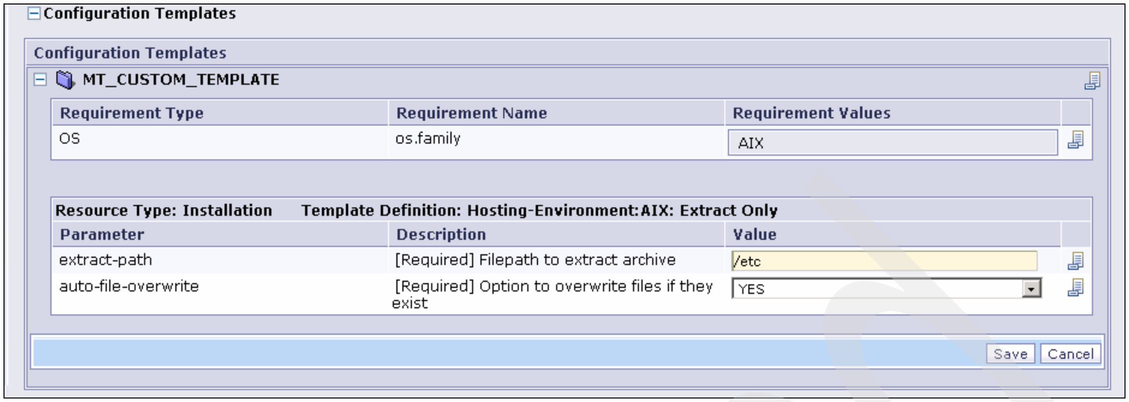

Configuration Templates

Configuration Templates

MT_CUSTOM_TEMPLATE

Requirement Type	Requirement Name	Requirement Values
OS	os.family	AIX

Resource Type: Installation   Template Definition: Hosting-Environment:AIX: Extract Only

Parameter	Description	Value
extract-path	[Required] Filepath to extract archive	/etc
auto-file-overwrite	[Required] Option to overwrite files if they exist	YES

Figure 3-78 Setting template parameter for installation

Click **Save**. We can now install the tar file on the target computer. Select **Edit** → **Install** as shown in Figure 3-79.

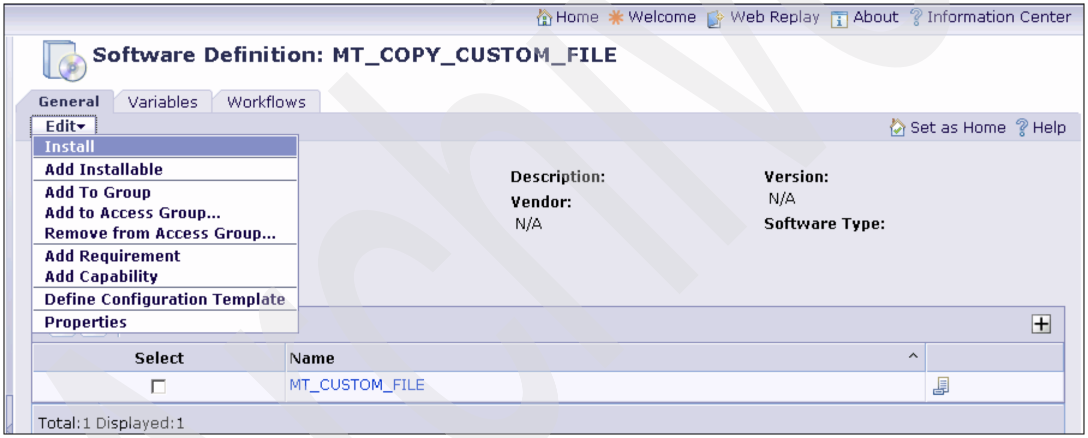

Software Definition: MT_COPY_CUSTOM_FILE

Description:

Vendor:
N/A

Version:
N/A

Software Type:

Select	Name
□	MT_CUSTOM_FILE

Total:1 Displayed:1

Figure 3-79 Install software product you created for customization

Select the software product you created, and select the computer as shown in Figure 3-80 on page 163.

TPM then copies the `custom.tar` file in the `/etc` directory of the target computer. TPM extracts it automatically and overwrites the files on the target computer.

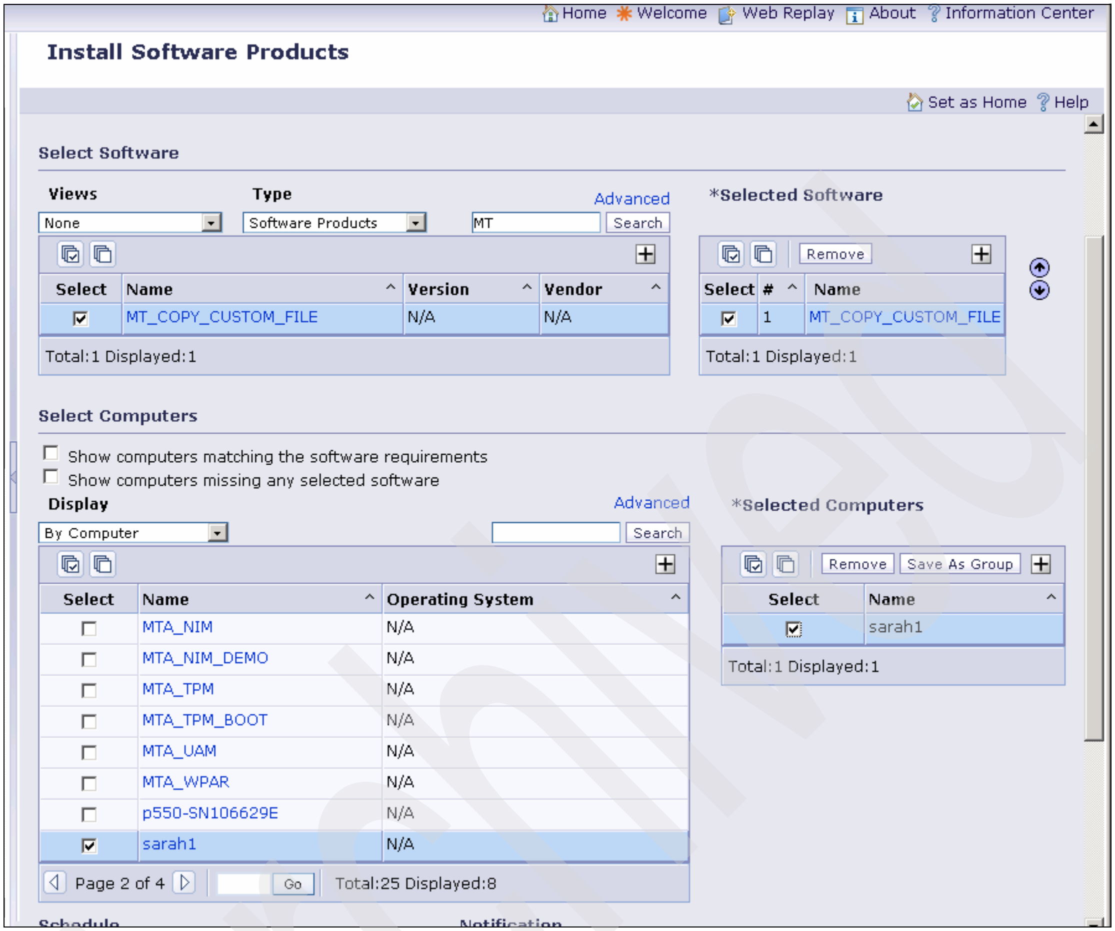

Home Welcome Web Replay About Information Center

Install Software Products

Set as Home Help

Select Software

Views: None | Type: Software Products | MT | Search | Advanced

Select	Name	Version	Vendor
☑	MT_COPY_CUSTOM_FILE	N/A	N/A

Total:1 Displayed:1

*Selected Software

Remove

Select	#	Name
☑	1	MT_COPY_CUSTOM_FILE

Total:1 Displayed:1

Select Computers

□ Show computers matching the software requirements
□ Show computers missing any selected software

Display: By Computer | Search | Advanced

Select	Name	Operating System
□	MTA_NIM	N/A
□	MTA_NIM_DEMO	N/A
□	MTA_TPM	N/A
□	MTA_TPM_BOOT	N/A
□	MTA_UAM	N/A
□	MTA_WPAR	N/A
□	p550-SN106629E	N/A
☑	sarah1	N/A

Page 2 of 4 | Go | Total:25 Displayed:8

*Selected Computers

Remove | Save As Group

Select	Name
☑	sarah1

Total:1 Displayed:1

Schedule | Notification

Figure 3-80 Install software product on specific virtual server

Operating system technology levels and bundles

To install a technology level (TL) or a service pack on AIX, TPM makes use of the Service Update Management Assistant (SUMA), which automates the update process for AIX systems.

SUMA for AIX, which is included with the AIX 5L™ Version 5.3 and higher, automates all manual tasks of retrieving maintenance updates from the IBM Support Web site. It automates tasks such as downloading the latest security updates or downloading an entire Technology Level (TL). For more information about SUMA, read about system management in *AIX 5L Differences Guide Version 5.3 Edition*, SG24-7463.

For automated patch management in the AIX environment, the following components are required (or optional where mentioned):

- Installed satellite server AIX, which can connect to AIX Fix Center through the Internet.
- FTP or HTTP proxy (optional)
- TPM server
- Installed endpoint servers AIX (Patches will be installed on these servers.)

The configuration is shown in Figure 3-81.

Using SUMA, the AIX satellite server connects to the IBM Fix Center server and performs the following operations:

- Scanning for available patches
- Downloading required patches for the endpoints
- Sending the recommended patches to the endpoints

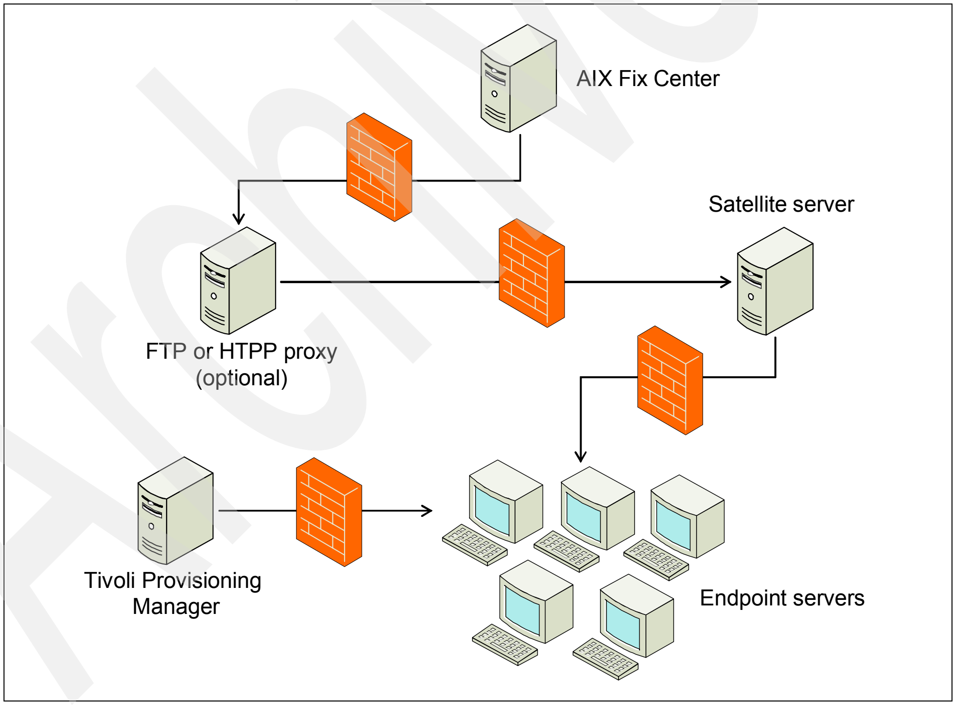

Figure 3-81   System configuration for patch management with TPM and SUMA

With this configuration, we implement patch management by TPM, as follows:

1. Discover satellite server by TPM and setup password credentials
2. Create discovery configuration. The type is `AIX Discover Patches`.
3. Scan endpoints to obtain missing patch status, and register status in inventory.
4. Install patches to endpoint servers. Use one of the following methods:
   - Install patches from recommendations
   - Install patches as a software stack
   - Install individual patches

Management product installation

After the server is installed and ready for service, we must install any necessary agents for the service provider to manage the server. Usually, these agents are common to all servers, and we can include the agent software when we make the base AIX image.

If you are using IBM Tivoli Monitoring, install Tivoli Enterprise Monitoring Agent in `/opt/IBM/TIM`. For more information, see 4.4.3, “IBM Tivoli Monitoring” on page 195.

If you are using Tivoli Usage and Accounting Manager (TUAM) and AIX Advanced accounting, install accounting agent in the following path:

```
/opt/ibm/tuam/collectors/unix
```

For more information, see 5.6, “Configuring data collection in AIX” on page 241.

User application installation

After the basic installation is complete, we can install appropriate user application software with TPM distribution infrastructure. TPM provides distribution and installation software products, software patches, and software stacks to target endpoint servers. With TPM, we can periodically scan the target endpoint servers to keep installed software inventory information up-to-date.

For application installation, use the following steps:

1. For data collection and management by TPM, install Tivoli Common Agent on managed endpoint server.
2. Publish the software products, software patches, or files in a specified repository to depot servers.
3. Make a schedule for installation.
4. Install the software as we discussed.

Installing operating system technology levels, bundles, and applications requires many steps. For specific implementation, search for managing operating systems and software, and managing patches in the Tivoli Provisioning Manager documentation:

http://publib.boulder.ibm.com/infocenter/tivihelp/v20r1/index.jsp

4

Monitoring resources in a multitenant environment

In this chapter, we discuss the monitoring of multitenant environments from various perspectives.

This chapter contains the following topics:

4.1 Monitoring a multitenant environment

As we discussed in Chapter 2, “Multitenancy, PowerVM, and partitions” on page 17, virtualization is one way of realizing multitenancy. However, it increases the infrastructure complexity and the number of components that should be monitored. Although we have multiple aspects of monitoring, this section discusses the monitoring of multitenant systems from the following aspects:

- Monitoring system availability

  Multitenant systems run workloads of many different customers, and the effect of a general system downtime can be enormous. In addition to making system components redundant, an important point is to keep track of the status of each physical or virtual component, to ensure the whole system runs smoothly. Quickly identifying a failure of either a physical or virtual component is also essential to shorten the downtime.

- Monitoring resource usage

  As we increase resource sharing, monitoring the resource usage is important so we can:

  - Identify the bottleneck of the monitored systems to avoid performance decrease
  - Detect or predict resource exhaustion
  - Detect spare resources for further resource utilization optimization

- Monitoring performance

  Performance monitoring has to be done from the user level, so it measures the end-to-end performance, instead of a component performance. For example, to determine whether Web site performance is acceptable, we have to simulate complete Web requests from different source networks, instead of monitoring the Web server process, CPU and memory usage, and number of Web requests. Because of the application-dependent nature of performance monitoring task, we do not discuss performance monitoring in this book.

4.2 Considerations of monitoring the multitenant environment

This section discusses the following considerations that are unique to monitoring the multitenant environment:

- Requirements for the tenant systems
- Monitoring network isolation
- Monitoring responsibilities

4.2.1 Requirements for the tenant systems

Certain monitoring tools require a specific agent or daemon to be running on the monitored servers. Also, these tools might require specific network configurations. In some instances, such as the example of the “Multitenancy case A” on page 33, the tenant has full control over the OS instance. A service provider should clearly define what are acceptable operations for a tenant to perform. Some operations, such as changing the host name or IP address, and stopping or uninstalling the monitoring software agent, can have an effect on the service provider monitoring infrastructure. These issues should be part of a service level agreement (SLA) or a document of understanding (DOU).

4.2.2 Monitoring network isolation

A typical monitoring system architecture consists of a monitoring server, monitored servers with a monitoring agent program installed, and the network that connects the elements. An example is shown in Figure 4-1 on page 170.

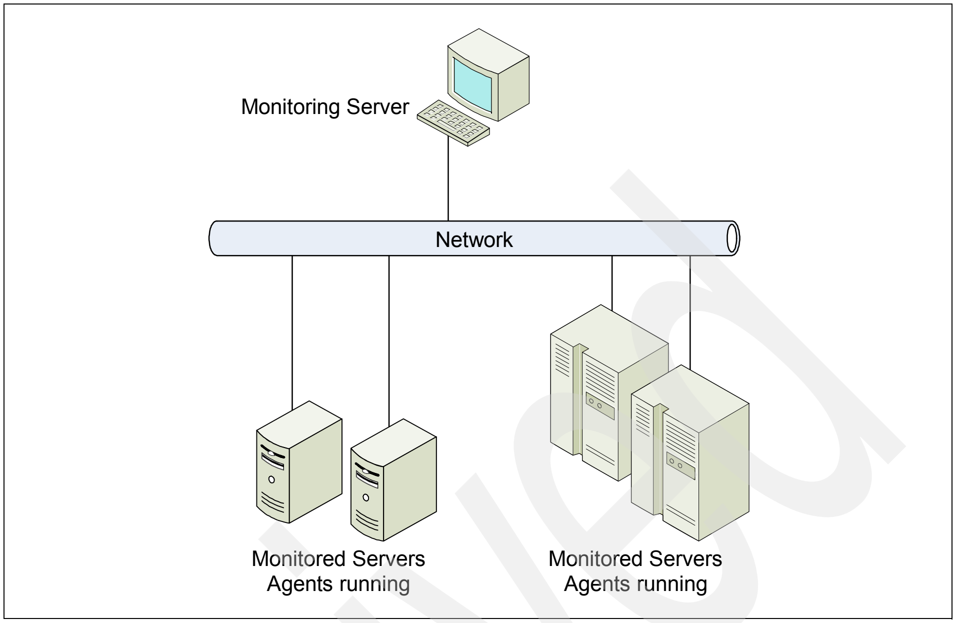

Figure 4-1 Monitoring system components example

Usually, monitoring and the production network are separated. The separation is often required because the networks require different security levels, and the monitoring itself can affect performance.

Figure 4-2 on page 171 below an example of a simple monitoring network configuration.

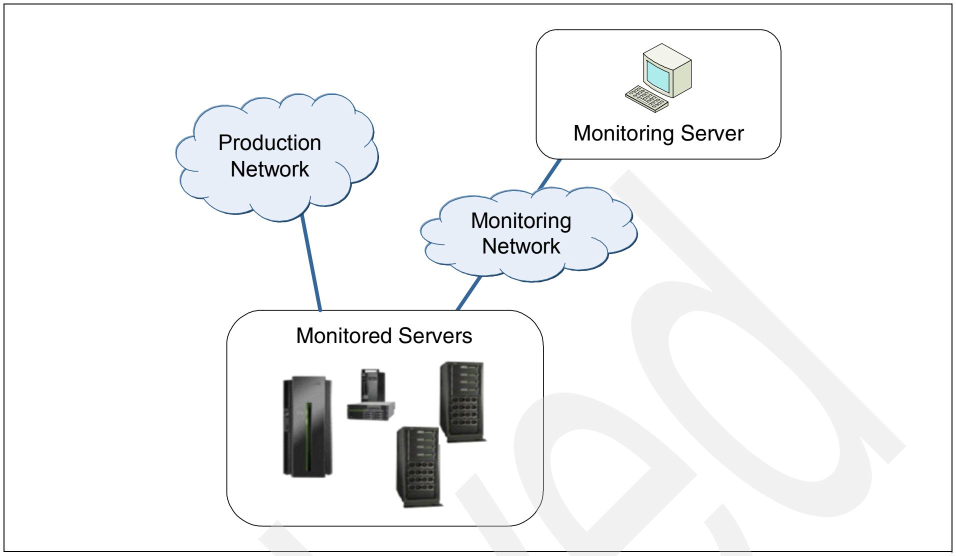

Figure 4-2 Simple monitoring system network

In a multitenant environment, in addition to separation of service and management, the isolation between tenants is also important. Your tenants do not want to be seen by other tenants through the monitoring network. To obtain the maximum degree of isolation between tenants, you may build discrete networks and monitoring systems for each tenant. However, from a service provider perspective, a totally separated monitoring architecture tends to be costly, and can lack the global system monitoring aspect. To provide both the required isolation and the integrated monitoring system management, we should have properly configured firewalls between each tenant and our monitoring servers.

Figure 4-3 on page 172 shows an example of an multitenant environment, with different tenants and separated networks.

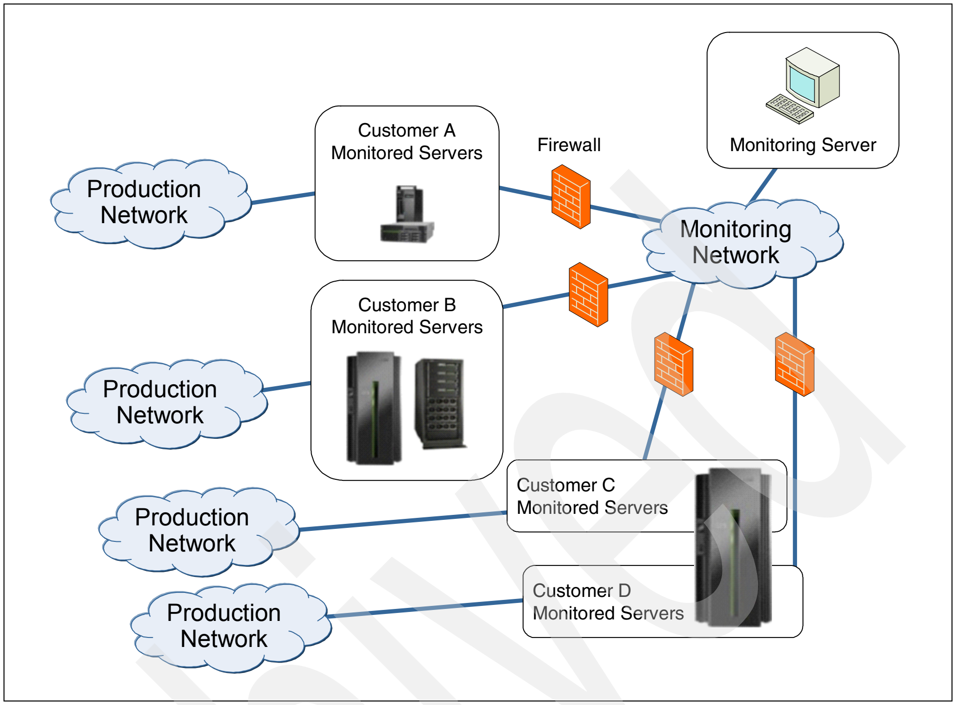

Figure 4-3 Multitenant monitoring system network

Every monitoring tool has different network requirements. Understand these requirements before you configure firewalls between tenants according to your security policies. We discuss the network requirements when we discuss each monitoring tool.

4.2.3 Monitoring responsibilities

Another consideration is to define who is responsible to monitor the physical components, virtual resources, or the applications in multitenant environment.

When monitoring shared and virtualized environments, and hosting multiple tenants, determine (for each system component) whether the service provider or the tenant is responsible for monitoring the status of the component. Such responsibility separation has to be defined for each tenant because each tenant might have a different SLA.

We assign the system resources to be monitored to either of two roles:

- Service provider administrator
- Tenant administrator

Table 4-1 lists the different administrators and the responsibilities, for each required status.

Table 4-1 Administrator and responsibilities

Administrator	Availability status	Resource utilization status
Service provider administrator	Responsible to monitor the physical and virtual layers, ensuring the availability of the system.	Must monitor several components to gather resource utilization data to ensure that enough resources are available in the system.
Tenant administrator	Responsible to monitor the tenant's own applications, operating system and private resources, such as external storage. The tenant does not have to monitor the physical components. The tenant might not even be aware of the hardware being used.	Tenants can use their own tools to monitor the private resources if they need to gather the resource utilization data. They do not have to monitor other resources.

In the table, we organized the components under the physical and virtual layers, divided the responsibilities, and defined the objectives when we performed monitoring.

Figure 4-4 on page 174 shows components that each administrator can monitor for availability information.

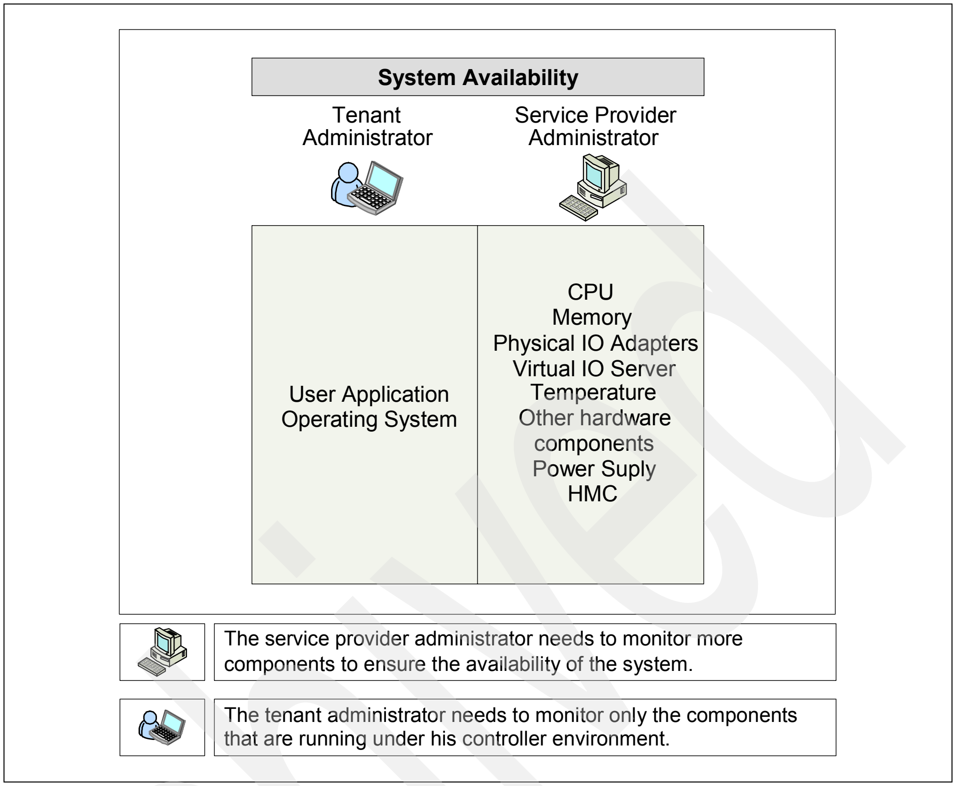

The service provider administrator needs to monitor more components to ensure the availability of the system.

The tenant administrator needs to monitor only the components that are running under his controller environment.

Figure 4-4 Responsibilities for monitoring system availability

Figure 4-5 on page 175 shows the resource utilization components and the different responsibilities.

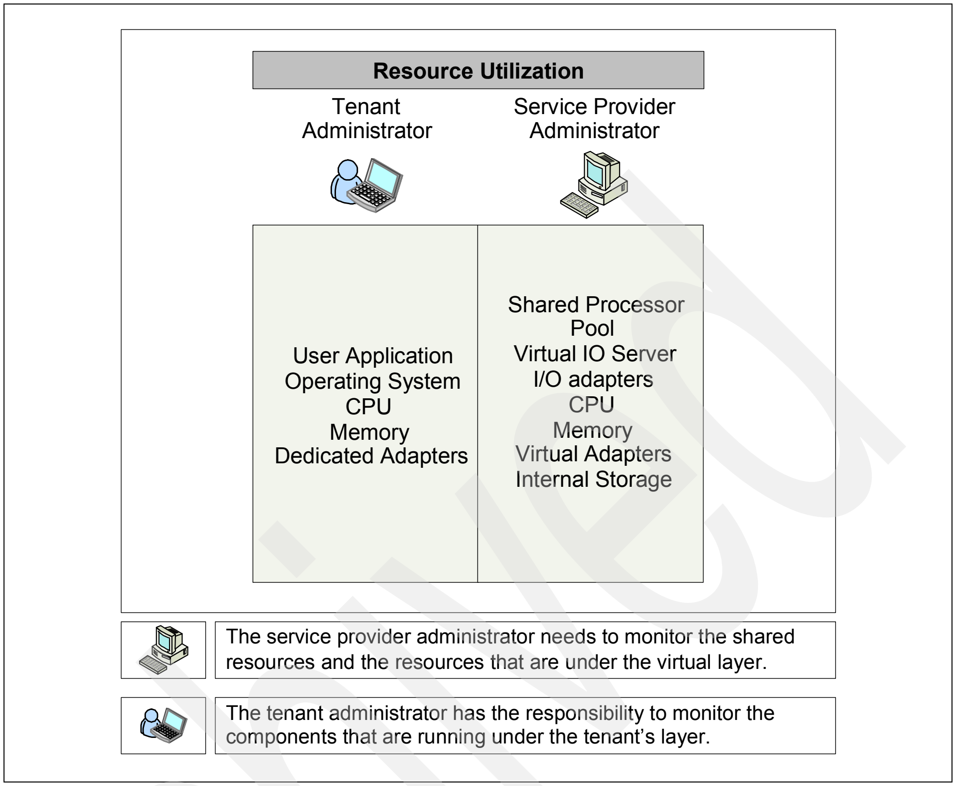
Resource Utilization
Tenant Administrator
Service Provider Administrator
User Application
Operating System
CPU
Memory
Dedicated Adapters
Shared Processor Pool
Virtual IO Server
I/O adapters
CPU
Memory
Virtual Adapters
Internal Storage
The service provider administrator needs to monitor the shared resources and the resources that are under the virtual layer.
The tenant administrator has the responsibility to monitor the components that are running under the tenant's layer.

Figure 4-5 Responsibilities for monitoring resource utilization

4.3 What to monitor in a multitenant environment

One of the concerns when monitoring a highly virtualized environment is how to identify the components that have to be monitored.

To identify the monitored components in multitenant Power Systems, we show the current IBM offerings for Power Systems technology, their major components, and how the components are divided into layers.

4.3.1 Monitored layers

We use the following terms to define the layers:

- Physical infrastructure layer

  Any Power Systems server including Hardware Management Console (HMC), hardware components such as CPU, memory, physical adapters and other hardware parts
- Virtual infrastructure layer

  Any PowerVM features such as logical partition (LPAR), Virtual I/O Server, and WPAR
- Tenant or application layer

  The tenant or application layer includes the operating system and the tenant's applications.

Physical infrastructure layer

This section discusses the available Power Systems servers and physical components (hardware) that belong to the physical infrastructure.

Power Systems servers

Figure 4-6 on page 177 shows the Power Systems servers used in our physical infrastructure.

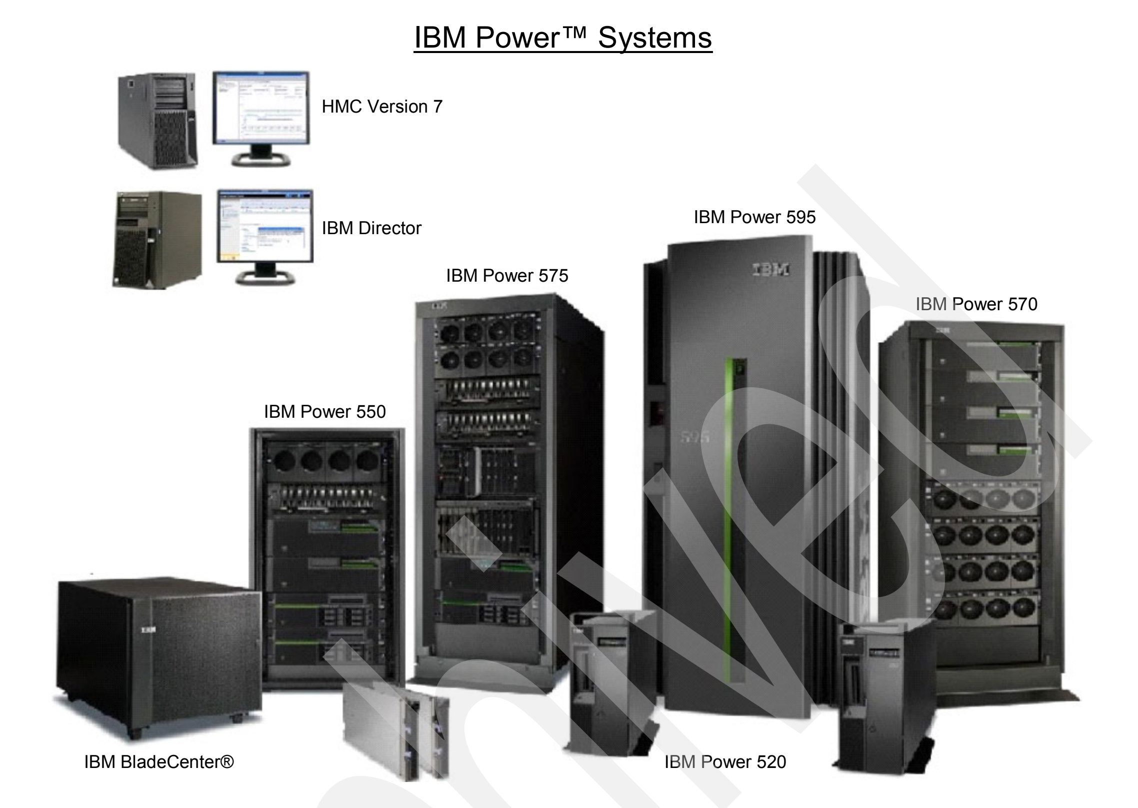

Figure 4-6 Power Systems servers currently available

Power Systems physical components

Each server has its own components, which the service provider administrator must monitor. Figure 4-7 on page 178 shows the physical infrastructure components and indicators.

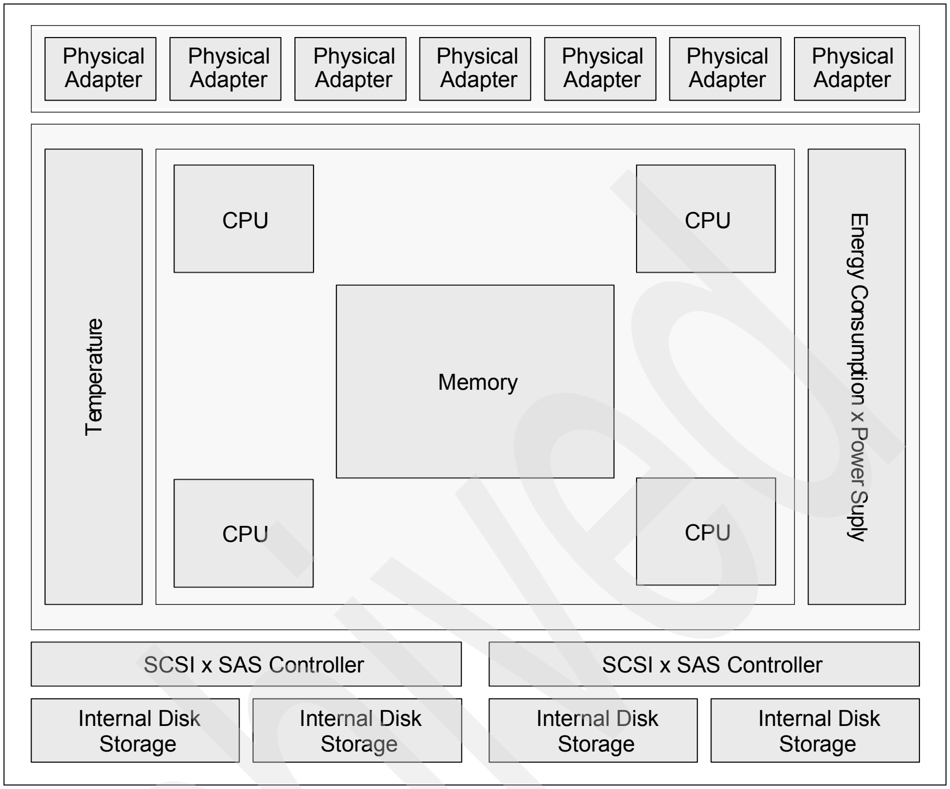

Figure 4-7 View of the major physical infrastructure components

For details about Power Systems hardware information, see the facts and features Web page:

http://www.ibm.com/systems/power/hardware/reports/factsfeatures.html

Virtual Infrastructure layer

The PowerVM technology has features to create virtual devices. Each virtual device can be used as a single device for a particular LPAR or can be shared among multiple clients. For this reason, the virtual components must be monitored. Figure 4-8 on page 179 shows an example of a virtualized configuration.

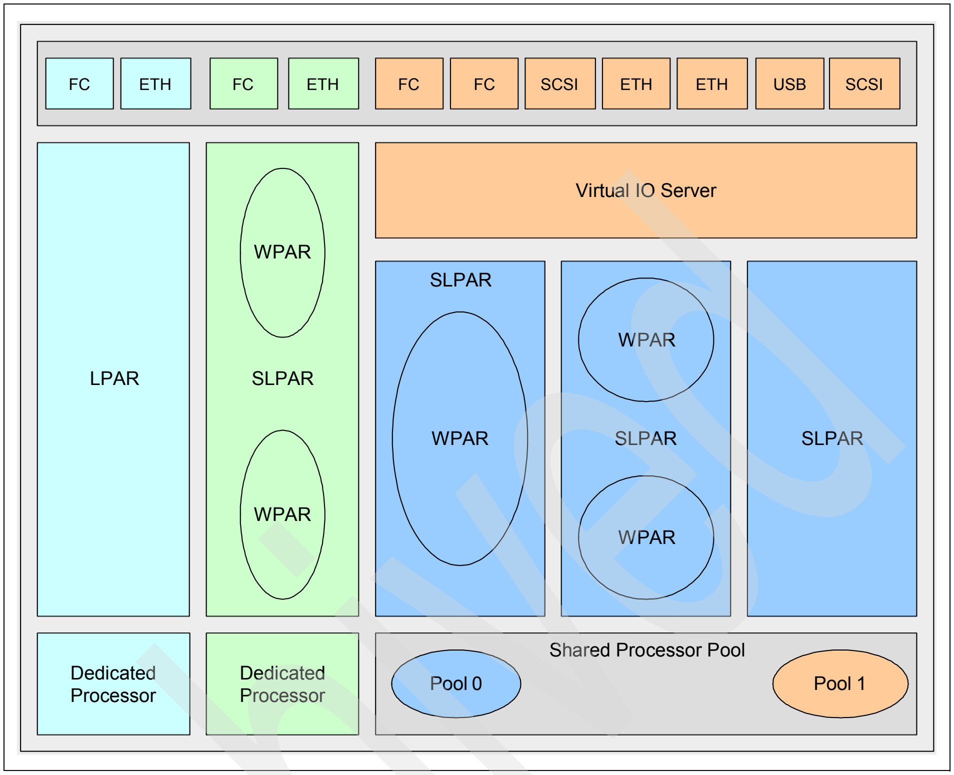

Figure 4-8 Virtual configuration using PowerVM features

For technical descriptions about PowerVM features, see 2.1, “PowerVM Virtualization technologies” on page 18.

Tenant or application layer

The tenant or application layer includes the operating system, various middleware tools, and the tenant’s applications.

4.3.2 Monitored aspects for each layer

This section discusses the monitored aspects for the physical, virtual, and tenant or application layers.

System availability

One of the objectives of the monitoring process, the system availability indicator, helps the multitenant administrator ensure the health of the system, keep the environment running, and perform maintenance when necessary.

This section shows the different layers and the components monitored to provide availability information. We also describe several of the failure types.

Physical infrastructure layer

The physical Infrastructure layer represents all the hardware components and indicators that the service provider administrator can monitor to provide availability status. We also mention the major types of failures that have to be monitored to ensure the availability of the environment.

- Hardware malfunctions

  A hardware malfunction can occur because of a physical problem, followed by a component failure. The following physical components can be affected:

  - CPU
  - Memory
  - Physical adapter
  - HMC
  - Internal storage
  - Power Systems server
- Environmental problems

  An environmental problem can occur from power or thermal problems. The following components or indicators can be monitored:

  - Energy problems
  - Temperature problems
- Network issues

  The network has a big effect in the availability of the environment. The most common problems are related to the network link. The following components can be monitored:

  - Physical network cards
  - Cables and network switches

Virtual infrastructure layer

The virtual infrastructure layer represents all virtual devices that are created by the PowerVM features. Several components that can be affected during a failure are listed in this section. Virtual components can have failures that are caused by a defect, in any of the components, or a hardware problem. The following components can provide availability information:

- Virtual I/O Server
- Virtual Ethernet problems
- Logical partition problems
- Virtual adapters

Tenant layer

The tenant layer contains the tenant's application or operating system that has to be monitored to ensure the availability of the system.

In the tenant layer, we can monitor the following components to provide availability status:

- Logical partition (LPAR)
- Workload partition (WPAR)
- User's application

The tenant administrator does not usually monitor this layer because it is maintained by the tenant.

System resource utilization

This section discusses the physical and virtual layers, and the components that can be monitored to provide resource utilization data.

Physical infrastructure layer

Certain physical components can be monitored to provide overall resource utilization, to help the service provider identify performance problems, or to ensure the system has all the resources correctly allocated.

We can monitor the following components in this layer:

- Shared processor pool

  The shared processor pool can be monitored by the multitenant service provider to provide overall resource utilization of the pool and to show which shared LPAR is consuming the most CPU cycles.

- Memory

  The multitenant service provider can monitor the usage of the memory within a managed server to keep track of allocated resources.

- Internal storage

  The internal storage in a multitenant environment is normally allocated to the Virtual I/O Server to provide access through virtual devices. It can be monitored to provide usage information and response time.

- Physical adapters

  The physical adapters can be monitored by the service provider administrator if they are used by the Virtual I/O Server. If the adapters are dedicated to the tenant system, they are usually monitored by the tenant.

Virtual infrastructure layer

The virtual components that can be monitored to provide utilization and performance metrics are:

- Virtual I/O Server

  The Virtual I/O Server can host several networks or virtual SCSI devices and it can provide CPU, memory, and I/O utilization to help the service provider administrator prevent possible bottlenecks.

- Virtual adapters

  Virtual adapters can provide bandwidth utilization to help the service provider administrator identify possible memory, CPU or virtual I/O problems.

Tenant or application layer

The tenant's environment can provide utilization information to help the tenant ensure the system is running properly. The following components can provide resource utilization data:

- Operating system

  The operating system can provide user and kernel performance data, and provide CPU and memory usage for each process and application.

- WPAR

  The WPAR can provide CPU and memory usage.

- User applications

  The user applications or programs can be monitored to provide process usage, user and profiling usage by CPU, memory, and I/O utilization.

4.4 Monitoring tools for multitenant systems

IBM offers multiple tools to support monitoring tasks. This section discusses the following tools, from the multitenancy on Power Systems perspective:

- Hardware Management Console
- Integrated Virtualization Manager
- IBM Tivoli Monitoring
- IBM Director
- AIX built-in commands
- Third-party commands

These monitoring tools can watch the availability and resource usage of whole systems or their components. Each tool is responsible for different characteristics and covers the components differently. Figure 4-9 illustrates the coverage of monitored tools.

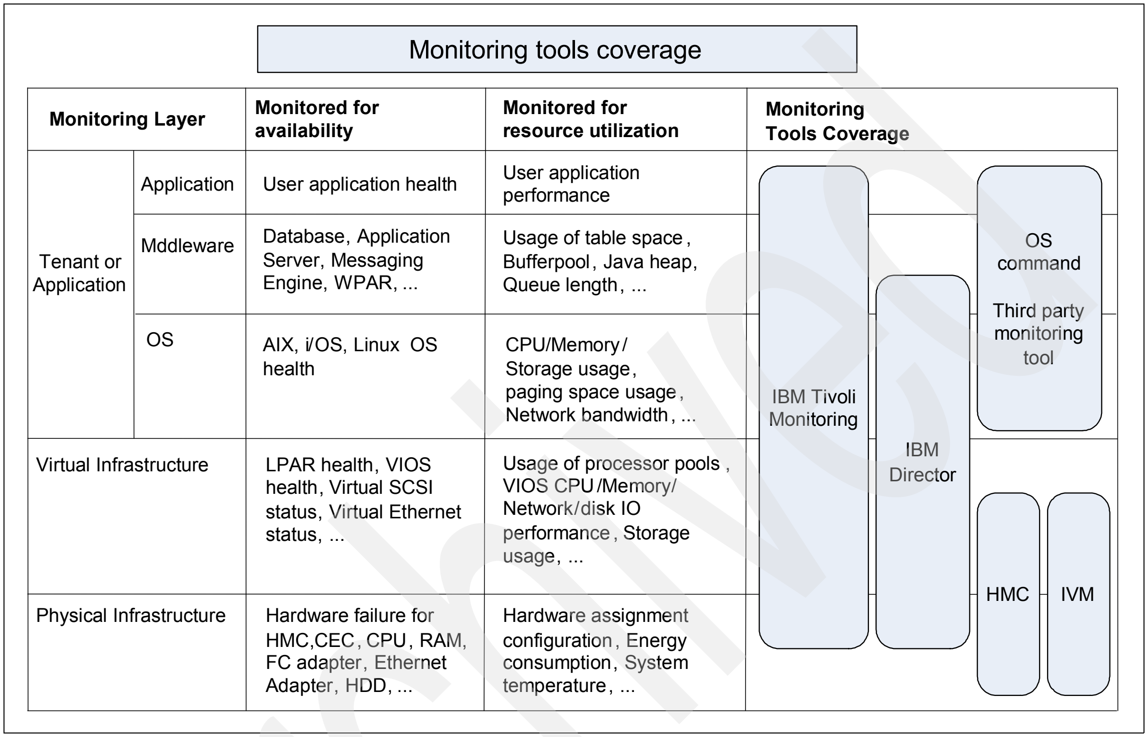

Monitoring tools coverage

Monitoring Layer		Monitored for availability	Monitored for resource utilization	Monitoring Tools Coverage
Tenant or Application	Application	User application health	User application performance	IBM Tivoli Monitoring OS command Third party monitoring tool
	Mddleware	Database, Application Server, Messaging Engine, WPAR, ...	Usage of table space, Bufferpool, Java heap, Queue length, ...	IBM Director
	OS	AIX, i/OS, Linux OS health	CPU/Memory/ Storage usage, paging space usage, Network bandwidth, ...	
Virtual Infrastructure		LPAR health, VIOS health, Virtual SCSI status, Virtual Ethernet status, ...	Usage of processor pools, VIOS CPU/Memory/ Network/disk IO performance, Storage usage, ...	HMC IVM
Physical Infrastructure		Hardware failure for HMC,CEC, CPU, RAM, FC adapter, Ethernet Adapter, HDD, ...	Hardware assignment configuration, Energy consumption, System temperature, ...	

Figure 4-9 Monitoring tools coverage outline

4.4.1 Hardware Management Console

This section applies to HMC Version 7, Release 3, Modification 30, Service Pack 2.

The IBM Hardware Management Console (HMC) is a management tool for Power Systems servers. System administrators use it for planning, deploying, and managing Power Systems servers. The major functions that the HMC provides are server hardware management and virtualization management, including system monitoring capabilities. The system administrator uses the functions through a Web browser or an Secure Shell (SSH) connection.

With HMC functions, system administrators can perform various management tasks for multiple servers from a single location.

Capabilities include:

- Creating, displaying, and maintaining LPARs
- Turning on or turning off the managed systems
- Opening a virtual terminal for each partition
- Displaying managed system resources and status
- Performing dynamic LPAR operations
- Managing Capacity on Demand operation
- Acting as a service forcal point

From a system monitoring perspective, HMC detects the hardware and LPAR failure events of managed systems, and reports the events to the predefined system administrators or service providers. HMC is also able to monitor and display the managed system's processor usage.

For details about HMC, see the following resources:

- *Hardware Management Console V7 Handbook*, SG24-7491
- Support and downloads for HMC

  `https://www14.software.ibm.com/webapp/set2/sas/f/hmc/resources.html#v7`
- AIX, VIOS and HMC Facts and Features

  `http://www.ibm.com/support/docview.wss?uid=tss1td103130&aid=1`

Monitoring features of HMC

HMC provides two types of function for monitoring:

- Detecting and reporting system failures
- Collecting processor utilization data

HMC detects hardware failures of managed systems. HMC can automatically report the failure information to specified destinations by e-mail or SNMP trap. If the call-home feature is set up, HMC also sends the information to a service provider, through phone line, Secure Sockets Layer (SSL), or virtual private network (VPN).

The other available information is the utilization of dedicated or shared processors. This can be useful to understand the actual usage of overall system processor usage, from the global system or LPAR perspective.

Detecting and reporting system failures

When you access HMC through a Web browser, the welcome Web page shows the summarized information about managed systems, as shown in Figure 4-10. This information appears before you log on to HMC. To obtain detailed information about managed systems, log on to HMC with a user ID that has certain roles assigned.

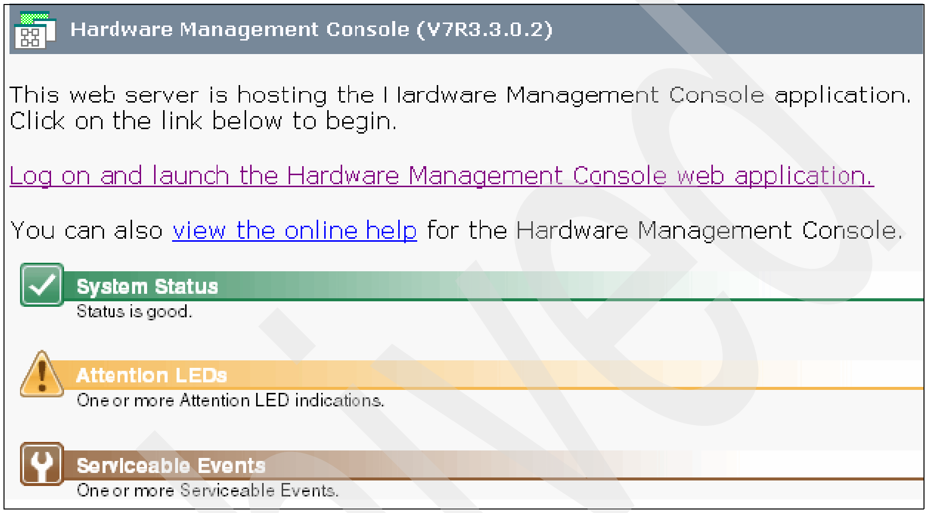

Hardware Management Console (V7R3.3.0.2)

This web server is hosting the I lardware Management Console application. Click on the link below to begin.

Log on and launch the Hardware Management Console web application.

You can also view the online help for the Hardware Management Console.

System Status
Status is good.

Attention LEDs
One or more Attention LED indications.

Serviceable Events
One or more Serviceable Events.

Figure 4-10 Welcome page of HMC

After logging into HMC, you can view even more detailed information by clicking the button at the bottom left corner of the HMC window, shown in Figure 4-11 on page 186. This information shows that the Attention LEDs are at the server level and indicate that the managed systems require attention, although are not necessarily critical. Both managed systems and LPARs have their own Attention LEDs.

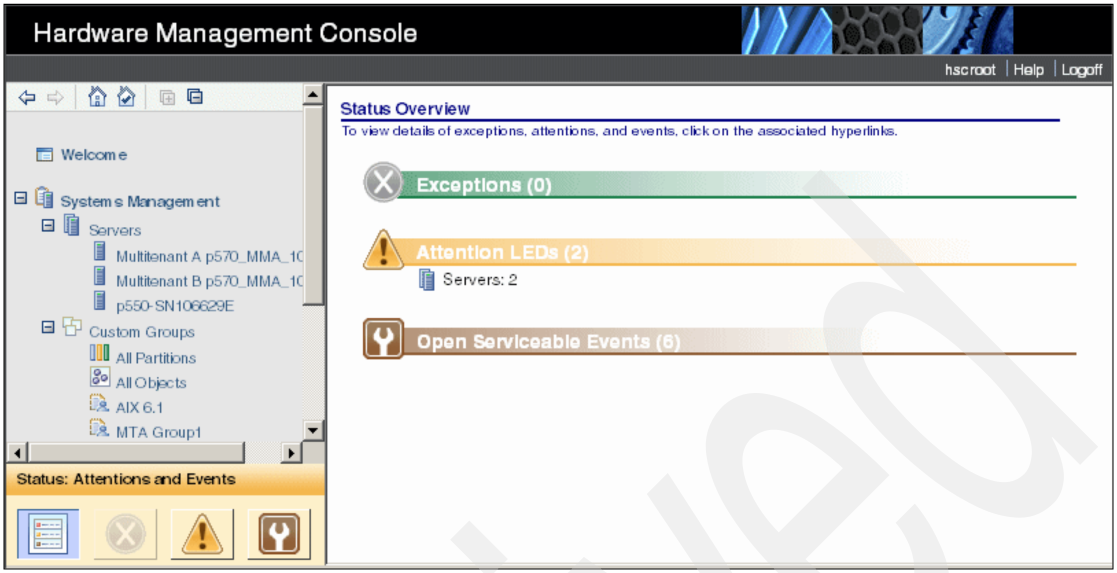

Hardware Management Console

hscroot | Help | Logoff

- Welcome
- ⊟ Systems Management
  - ⊟ Servers
    - Multitenant A p570_MMA_10
    - Multitenant B p570_MMA_10
    - p550-SN106629E
  - ⊟ Custom Groups
    - All Partitions
    - All Objects
    - AIX 6.1
    - MTA Group1

Status: Attentions and Events

Status Overview

To view details of exceptions, attentions, and events, click on the associated hyperlinks.

Exceptions (0)

Attention LEDs (2)

Servers: 2

Open Serviceable Events (6)

Figure 4-11 Managed systems status overview

If any Attention LEDs are on, determine which server or LPAR is signalling the problem. To view the detailed LED status, select **Servers** in the left pane, as show in Figure 4-12. Normally these LEDs are activated when serviceable events are triggered.

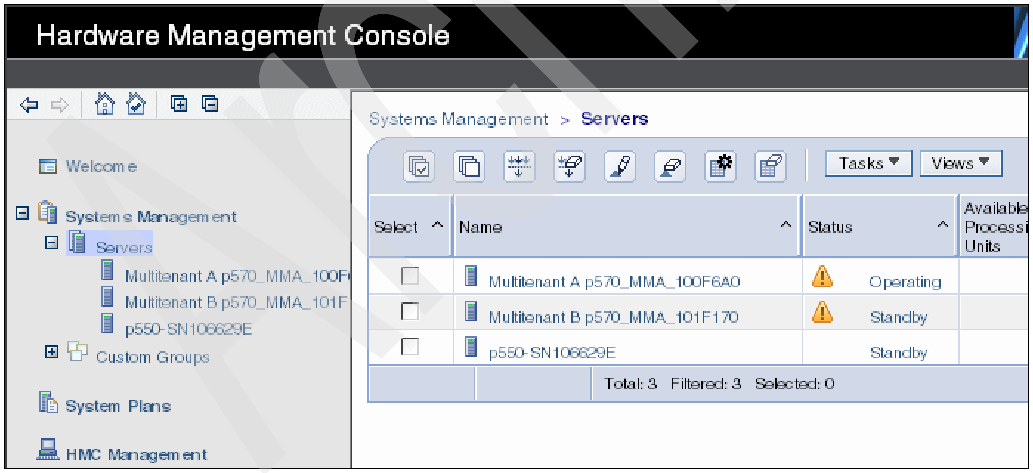

Hardware Management Console

- Welcome
- ⊟ Systems Management
  - ⊟ Servers
    - Multitenant A p570_MMA_100F
    - Multitenant B p570_MMA_101F
    - p550-SN106629E
  - ⊞ Custom Groups
- System Plans
- HMC Management

Systems Management > **Servers**

Select ^	Name ^	Status ^	Available Processi Units
□	Multitenant A p570_MMA_100F6A0	⚠ Operating	
□	Multitenant B p570_MMA_101F170	⚠ Standby	
□	p550-SN106629E	Standby	

Total: 3 Filtered: 3 Selected: 0

Figure 4-12 Managed systems attention LED status

To view additional details about managed systems events, open the **Serviceable Events Overview** window as shown in Figure 4-14 on page 188. To open the window, click **Open Serviceable Events** shown in Figure 4-11 on page 186.

To filter the events, selecting the criteria as shown in Figure 4-13.

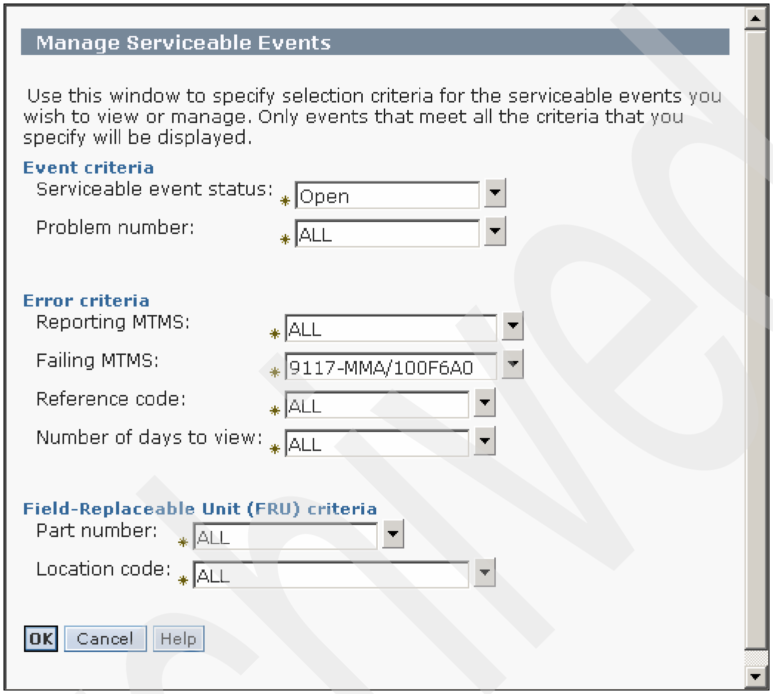

Manage Serviceable Events

Use this window to specify selection criteria for the serviceable events you wish to view or manage. Only events that meet all the criteria that you specify will be displayed.

Event criteria

Serviceable event status: * Open

Problem number: * ALL

Error criteria

Reporting MTMS: * ALL

Failing MTMS: * 9117-MMA/100F6A0

Reference code: * ALL

Number of days to view: * ALL

Field-Replaceable Unit (FRU) criteria

Part number: * ALL

Location code: * ALL

Figure 4-13 Setting criteria for serviceable events

An example list of serviceable events is shown in Figure 4-14 on page 188. To view detailed information, click a code in the Reference codes column. If the call-home function is set up correctly, HMC automatically calls the service provider, such as IBM, and sends the failure event information. This process helps to reduce the response time and to recover from the failures.

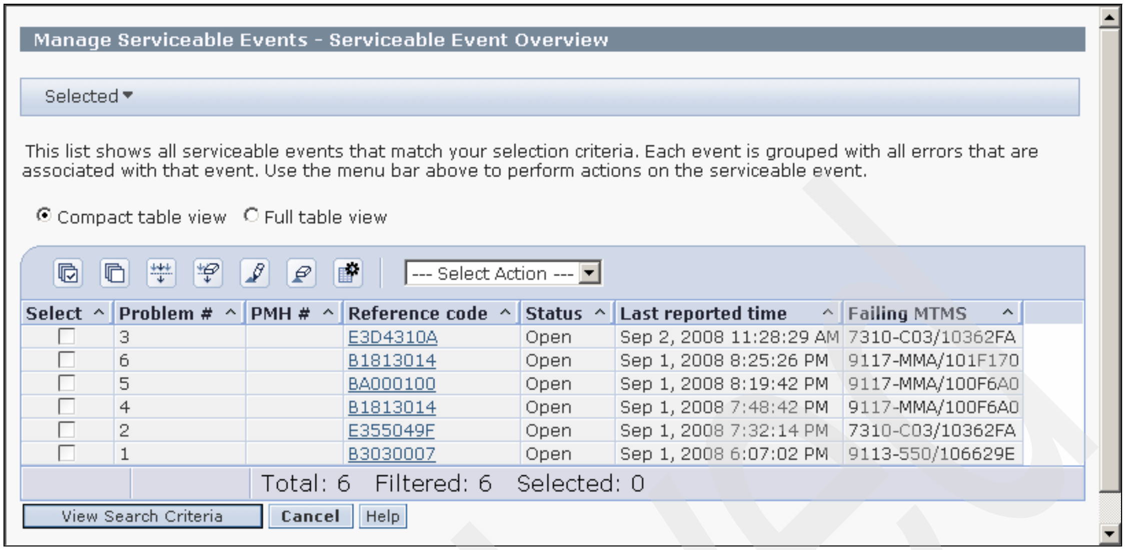

This list shows all serviceable events that match your selection criteria. Each event is grouped with all errors that are associated with that event. Use the menu bar above to perform actions on the serviceable event.

● Compact table view ○ Full table view

Select	Problem #	PMH #	Reference code	Status	Last reported time	Failing MTMS
□	3		E3D4310A	Open	Sep 2, 2008 11:28:29 AM	7310-C03/10362FA
□	6		B1813014	Open	Sep 1, 2008 8:25:26 PM	9117-MMA/101F170
□	5		BA000100	Open	Sep 1, 2008 8:19:42 PM	9117-MMA/100F6A0
□	4		B1813014	Open	Sep 1, 2008 7:48:42 PM	9117-MMA/100F6A0
□	2		E355049F	Open	Sep 1, 2008 7:32:14 PM	7310-C03/10362FA
□	1		B3030007	Open	Sep 1, 2008 6:07:02 PM	9113-550/106629E

Total: 6 Filtered: 6 Selected: 0

Figure 4-14 Serviceable events list example

Collecting processor utilization data

Processor utilization data is collected according to the specified retrieval rate. The default interval of the rate is 3600 seconds.

Overall processor utilization data is useful for multitenant system. With this information you can suggest the amount of CPU processing resources that is appropriate for tenants. It can also be used to plan upgrades of whole systems.

This processor utilization collection function does not require a network connection between the HMC and LPAR of the tenant.

To view the processor utilization data of a managed system (see Figure 4-15 on page 189):

1. Select **Servers** in the left pane.
2. Select the check box for the server (the figure shows that Multitenant A p570_MMA_100F6AO is selected).
3. Select **Operations** → **Utilization Data** → **View**.

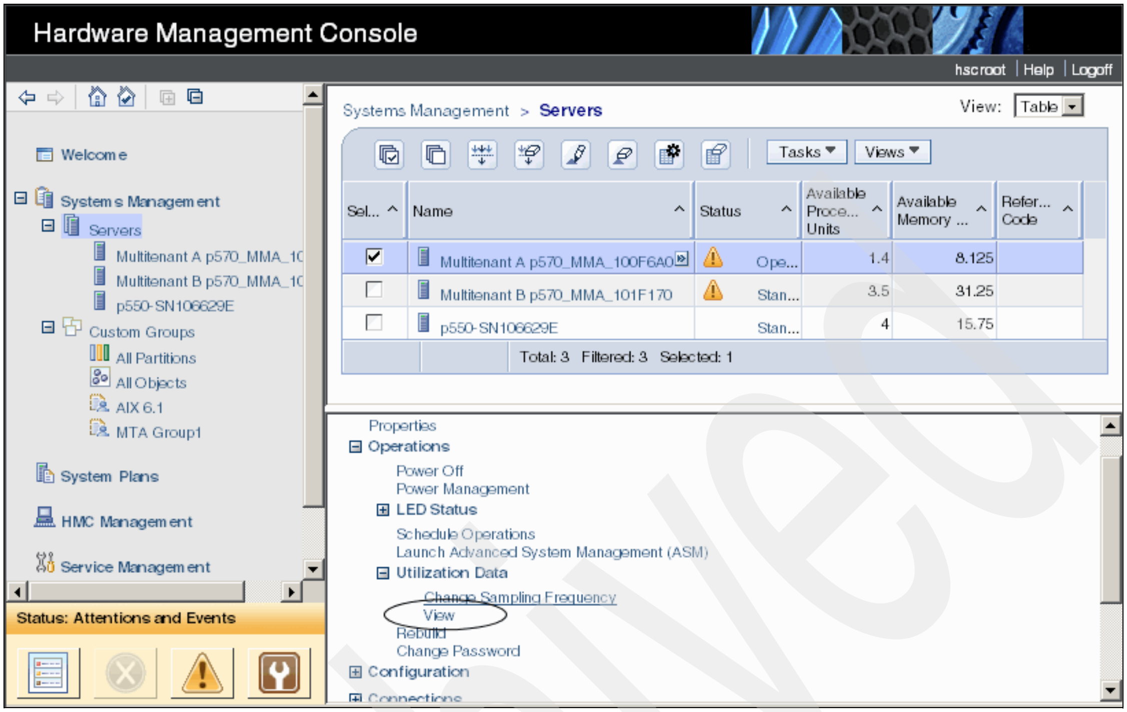

Sel...	Name	Status	Available Proce... Units	Available Memory ...	Refer... Code
☑	Multitenant A p570_MMA_100F6A0	Ope...	1.4	8.125	
☐	Multitenant B p570_MMA_101F170	Stan...	3.5	31.25	
☐	p550-SN106629E	Stan...	4	15.75	

Total: 3 Filtered: 3 Selected: 1

Figure 4-15 Displaying utilization data

To filter the results, select how often and specify a date range, as shown in Figure 4-16, and then click **Show**.

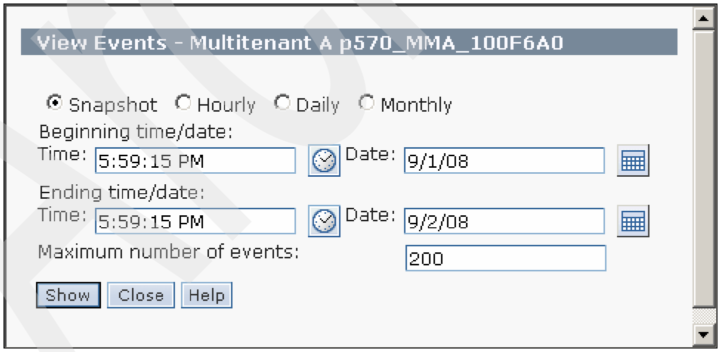
View Events - Multitenant A p570_MMA_100F6A0

● Snapshot ○ Hourly ○ Daily ○ Monthly

Beginning time/date:
Time: 5:59:15 PM Date: 9/1/08

Ending time/date:
Time: 5:59:15 PM Date: 9/2/08

Maximum number of events: 200

Figure 4-16 Filtering utilization data

The Utilization Events panel opens, as shown in Figure 4-17 on page 190. Select an event and click **Show Detail**. You cannot display multiple utilization data samples at once.

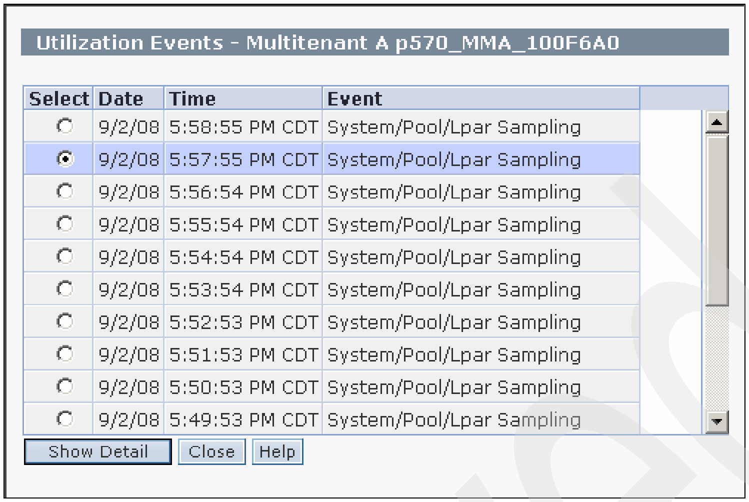

Utilization Events - Multitenant A p570_MMA_100F6A0

Select	Date	Time	Event
○	9/2/08	5:58:55 PM CDT	System/Pool/Lpar Sampling
●	9/2/08	5:57:55 PM CDT	System/Pool/Lpar Sampling
○	9/2/08	5:56:54 PM CDT	System/Pool/Lpar Sampling
○	9/2/08	5:55:54 PM CDT	System/Pool/Lpar Sampling
○	9/2/08	5:54:54 PM CDT	System/Pool/Lpar Sampling
○	9/2/08	5:53:54 PM CDT	System/Pool/Lpar Sampling
○	9/2/08	5:52:53 PM CDT	System/Pool/Lpar Sampling
○	9/2/08	5:51:53 PM CDT	System/Pool/Lpar Sampling
○	9/2/08	5:50:53 PM CDT	System/Pool/Lpar Sampling
○	9/2/08	5:49:53 PM CDT	System/Pool/Lpar Sampling

Figure 4-17 Selecting utilization event time stamp

System resource summary is displayed as shown in Figure 4-18. You can select the detailed information, such as LPAR, Physical Processor Pool, or Shared Processor Pool on the View pull-down menu.

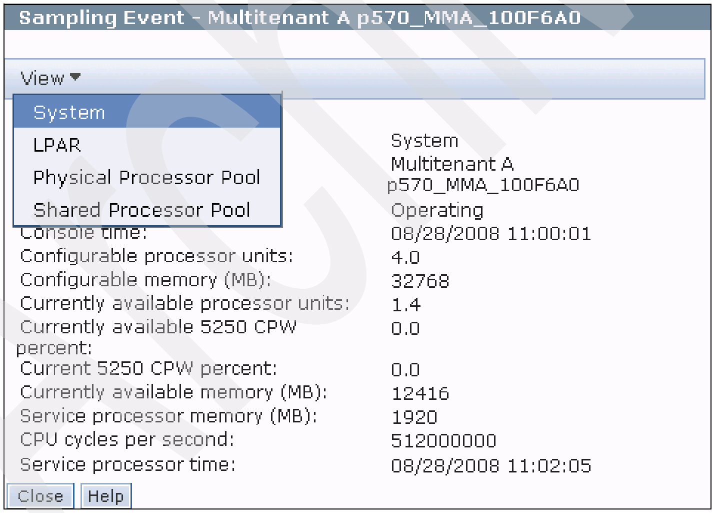

Sampling Event - Multitenant A p570_MMA_100F6A0

	System
	Multitenant A
	p570_MMA_100F6A0
	Operating
Console time:	08/28/2008 11:00:01
Configurable processor units:	4.0
Configurable memory (MB):	32768
Currently available processor units:	1.4
Currently available 5250 CPW percent:	0.0
Current 5250 CPW percent:	0.0
Currently available memory (MB):	12416
Service processor memory (MB):	1920
CPU cycles per second:	512000000
Service processor time:	08/28/2008 11:02:05

Figure 4-18 Specifying a view for utilization data

Figure 4-19 on page 191 is an example of the Physical Processor Pool view.

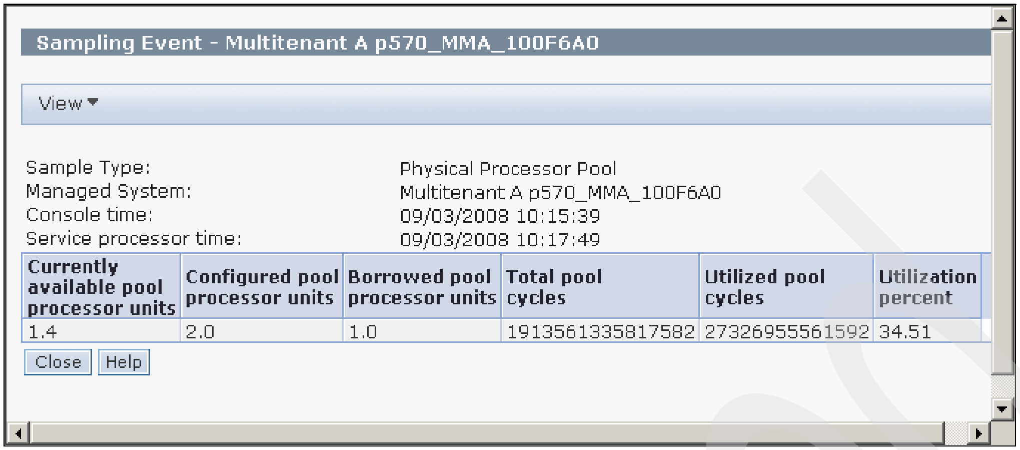
Sampling Event - Multitenant A p570_MMA_100F6A0

Sample Type:	Physical Processor Pool
Managed System:	Multitenant A p570_MMA_100F6A0
Console time:	09/03/2008 10:15:39
Service processor time:	09/03/2008 10:17:49

Currently available pool processor units	Configured pool processor units	Borrowed pool processor units	Total pool cycles	Utilized pool cycles	Utilization percent
1.4	2.0	1.0	1913561335817582	27326955561592	34.51

Figure 4-19 Physical Processor Pool view

Figure 4-20 is an example of the Shared Processor Pool view.

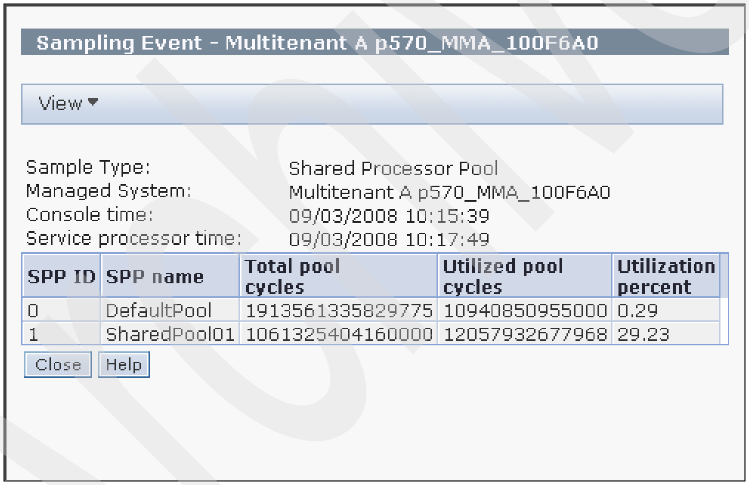
Sampling Event - Multitenant A p570_MMA_100F6A0

Sample Type:	Shared Processor Pool
Managed System:	Multitenant A p570_MMA_100F6A0
Console time:	09/03/2008 10:15:39
Service processor time:	09/03/2008 10:17:49

SPP ID	SPP name	Total pool cycles	Utilized pool cycles	Utilization percent
0	DefaultPool	1913561335829775	10940850955000	0.29
1	SharedPool01	1061325404160000	12057932677968	29.23

Figure 4-20 Shared Processor view

An example of the LPAR view is shown in two parts in Figure 4-21 and Figure 4-22.

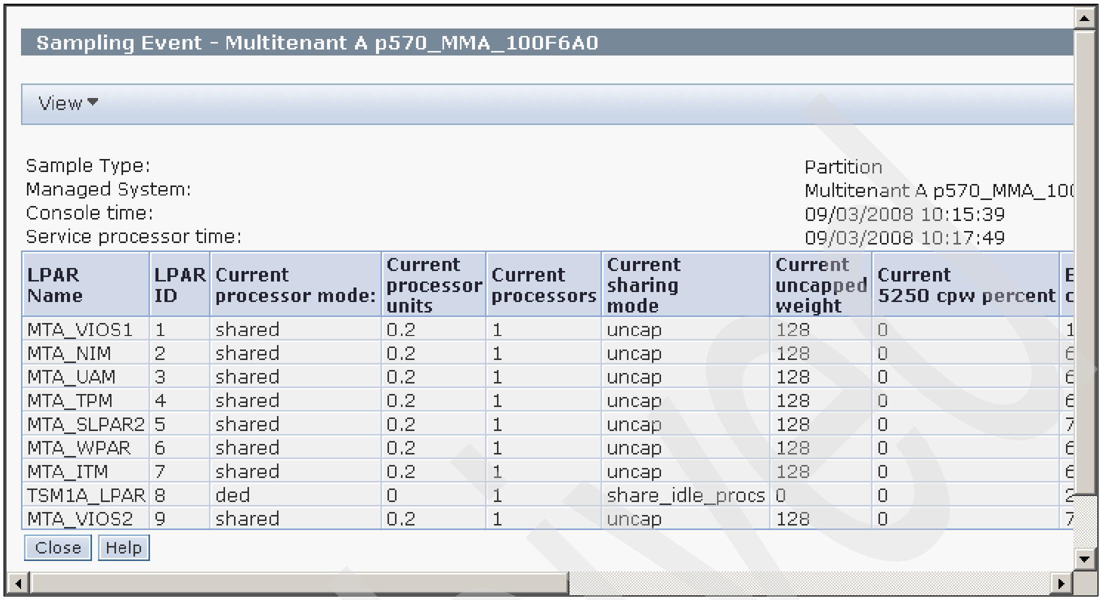

Sampling Event - Multitenant A p570_MMA_100F6A0

Sample Type: Partition
Managed System: Multitenant A p570_MMA_10(
Console time: 09/03/2008 10:15:39
Service processor time: 09/03/2008 10:17:49

LPAR Name	LPAR ID	Current processor mode:	Current processor units	Current processors	Current sharing mode	Current uncapped weight	Current 5250 cpw percent	E c
MTA_VIOS1	1	shared	0.2	1	uncap	128	0	1
MTA_NIM	2	shared	0.2	1	uncap	128	0	6
MTA_UAM	3	shared	0.2	1	uncap	128	0	6
MTA_TPM	4	shared	0.2	1	uncap	128	0	6
MTA_SLPAR2	5	shared	0.2	1	uncap	128	0	7
MTA_WPAR	6	shared	0.2	1	uncap	128	0	6
MTA_ITM	7	shared	0.2	1	uncap	128	0	6
TSM1A_LPAR	8	ded	0	1	share_idle_procs	0	0	2
MTA_VIOS2	9	shared	0.2	1	uncap	128	0	7

Figure 4-21 LPAR view (part 1)

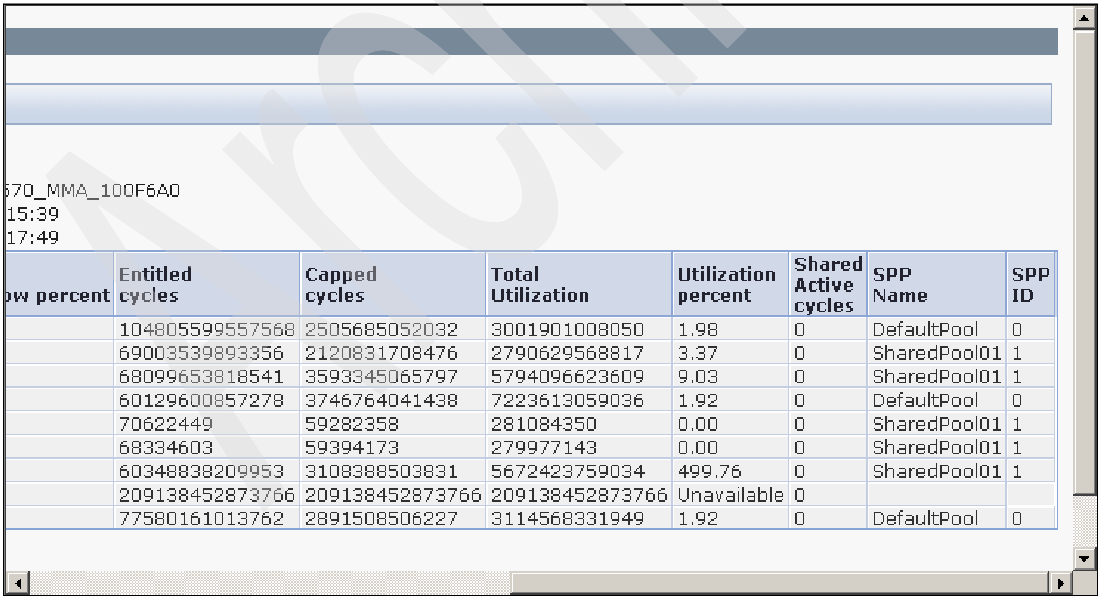

570_MMA_100F6A0
15:39
17:49

w percent	Entitled cycles	Capped cycles	Total Utilization	Utilization percent	Shared Active cycles	SPP Name	SPP ID
	104805599557568	2505685052032	3001901008050	1.98	0	DefaultPool	0
	69003539893356	2120831708476	2790629568817	3.37	0	SharedPool01	1
	68099653818541	3593345065797	5794096623609	9.03	0	SharedPool01	1
	60129600857278	3746764041438	7223613059036	1.92	0	DefaultPool	0
	70622449	59282358	281084350	0.00	0	SharedPool01	1
	68334603	59394173	279977143	0.00	0	SharedPool01	1
	60348838209953	3108388503831	5672423759034	499.76	0	SharedPool01	1
	209138452873766	209138452873766	209138452873766	Unavailable	0		
	77580161013762	2891508506227	3114568331949	1.92	0	DefaultPool	0

Figure 4-22 LPAR view (part 2)

HMC firewall considerations

HMC requires several different network communications as follows:

- HMC to managed systems connection

  This connection performs most of the hardware management functions in which HMC issues control function requests through the service processor of the managed system. This connection is mandatory for Power Systems server to be managed, except for blade servers.

- HMC to LPAR connection

  This connection collects platform-related information, such as hardware error events or hardware inventory, from the operating system that is running in the LPARs. It also coordinates certain platform activities, such dynamic LPAR or concurrent maintenance with those operating systems. This type of connection is optional.

- HMC to remote users connection

  This connection provides remote users with access to HMC functionality.

- HMC to service and support connection

  This connection transmits data such as hardware error reports, inventory data, and microcode updates, to and from your service provider. You can use this communication path to make automatic service calls.

Among these, the HMC to LPAR connection and HMC to remote users connection are connections between environment of tenant and HMC. Therefore, network isolation using firewalls between tenants has to be considered.

HMC to LPAR connection ports

HMC to LPAR connection uses the port shown in Table 4-2. This connection is required when using dynamic LPAR functionality. These ports must be opened on both HMC and LPAR because the connections are opened bidirectionally.

Table 4-2 Ports used for connections between HMC and LPAR

Service	Source	Destination	SRC port	DST port
Resource Monitoring and Control (RMC)	HMC	Monitored LPAR	Any	657:udp 657:tcp

HMC to remote users connection ports

HMC to remote users connection uses the port shown in Table 4-3 on page 194. For command line access, only the SSH port is required. For Web browser access, only the secure Web service ports have to be opened.

Table 4-3 Ports used by HMC client sessions

Service	Source	Destination	SRC port	DST port
Secure Web Service	Client PC	HMC	Any	443:tcp 8443:tcp 9660:tcp
SSH	Client PC	HMC	Any	22:tcp

For more information about ports used by HMC, see HMC Firewall Information:

http://www.ibm.com/support/docview.wss?uid=nas1387a6235643483f186256fee005d4c2c

4.4.2 Integrated Virtualization Manager

Integrated Virtualization Manager (IVM) provides a Web-based system management interface and a command-line interface that you can use to manage certain IBM Power Systems and IBM BladeCenter® blade servers. The IVM is included within the Virtual I/O Server, but it is activated and usable only on certain platforms. You can use IVM only when the managed Power Systems server is not managed by HMC.

IVM provides a subset of HMC functions, enabling you to perform tasks such as:

- Create and manage logical partitions
- Create and manage user accounts
- Create and manage serviceable events through IVM Service Focal Point™ feature
- Configure the virtual Ethernet networks
- Manage storage in the Virtual I/O Server
- Download and install updates to device microcode and to Virtual I/O Server software
- Back up and restore logical partition configuration information
- View application logs and the device inventory

System monitoring features of IVM

IVM can detect the managed system’s hardware malfunction and LPAR failure events, and it provides this information to a system administrator through the Service Focal Point. Unlike HMC, the Power Systems processor utilization data is not available in IVM.

Considerations with IVM

IVM has limited functions in comparison to HMC, including:

- No redundant Virtual I/O Server partition is supported.Only one Virtual I/O Server partition is allowed in a single Power System.
- Serviceability function is limited. No call-home functionality is available in IVM.

Because HMC offers better serviceability and capability of redundant configuration, consider using HMC instead of using IVM if you have both options available.

For detailed information about IVM, see:

- *PowerVM Virtualization on IBM System p: Managing and Monitoring,* SG24-7590
- CPWR6AA1 For Our Customers! Servicing the IBM System i® and System p Supporting POWER6 Technology. This course is available from:

  https://www.ibm.com/servers/resourcelink/lib03030.nsf/pagesbydocid/3bb5a3d1bdf144f6852571c5006014b4/$file/cover.htm

4.4.3 IBM Tivoli Monitoring

IBM Tivoli Monitoring (ITM)[1] monitors enterprise systems consisting of a wide variety of system components, from the perspective of system resource utilization and availability, and at multi-layered levels such as hardware, operating system, middleware, and application level. ITM provides a console to visualize, in real-time. It provides historical monitoring data for monitored systems in various formats such as charts and tables, which can help system administrators to monitor very complex systems easily. ITM also provides capability to automatically execute specified actions required for system problems. ITM is helpful for both service provider or tenant system administrators.

IBM Tivoli Monitoring components

Tivoli Monitoring consists of the components shown in Figure 4-23 on page 196. These components can be installed onto a single server for simplicity, or onto multiple servers for higher scalability and availability. The monitored system must have the agent module installed.

[1] This topic applies to IBM Tivoli Monitoring V6.2 Fix Pack level 1 interim fix level 1.

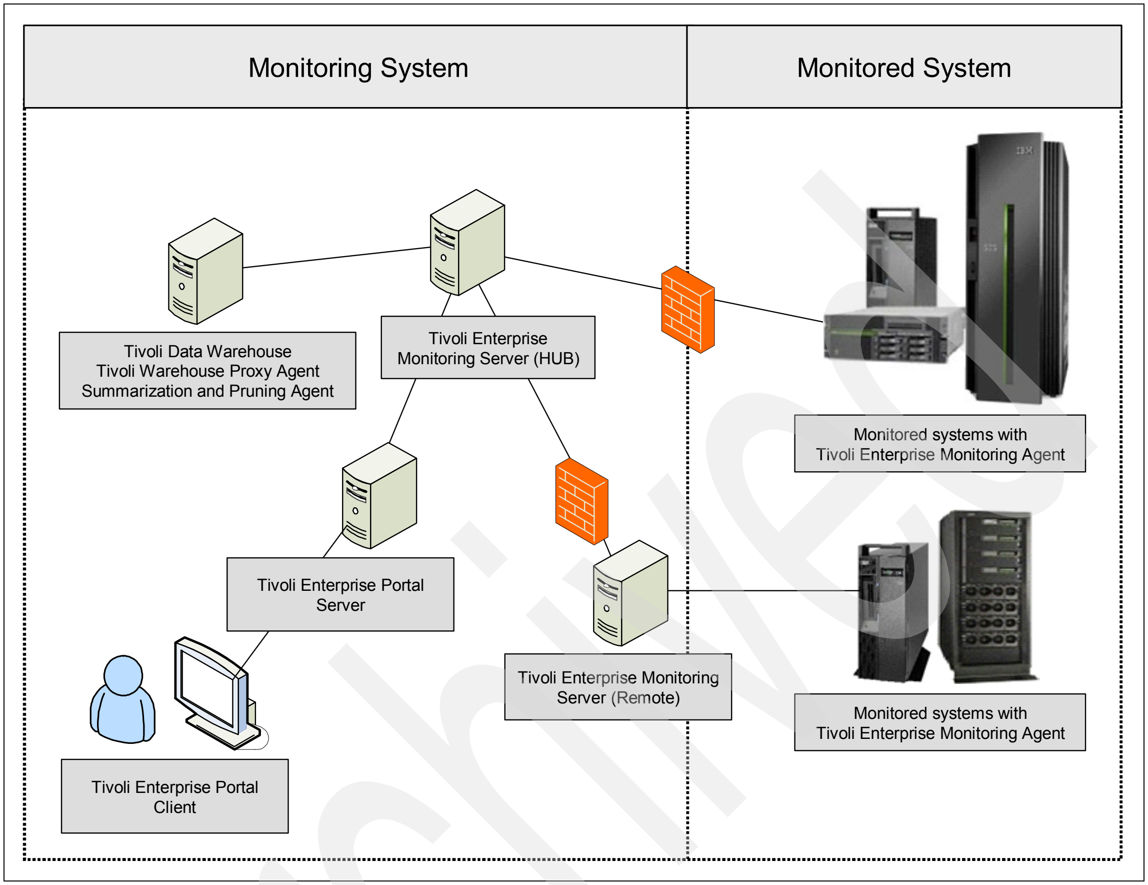

Figure 4-23 Tivoli Monitoring components overview

Tivoli Enterprise Monitoring Agents

The Tivoli Enterprise Monitoring Agents, also referred to as monitoring agents, are installed on the system or subsystem that requires data collection and monitoring. Tivoli Monitoring has various kinds of agents, each of which is responsible for collecting data for a specific component, such as an operating system, middleware, or an application.

Tivoli Enterprise Monitoring Agents are categorized into three groups:

- Operating system agents

  Operating system agents retrieve and collect all monitoring attribute groups related to specific operating system management conditions and associated data.

- Application agents

  These agents are specialized agents designed to retrieve and collect unique monitoring attribute groups related to one specific status of monitored component, such as databases, clustering tools, groupware tools, and so on.

- Universal agent

  This monitoring agent you can be configured to monitor any data you collect. It enables you to integrate data from virtually any platform and any source, such as custom applications, databases, systems, and subsystems.

Terms used in the Tivoli Monitoring environment are:

Attributes The monitored properties of systems collected by monitoring agents are called attributes. For example, CPU utilization, memory consumption, or number of running process are the attributes for operating system agents.

Situations Tivoli Monitoring tests the attributes collected by monitoring agents against predefined thresholds, and these tests are called situations**.** You can use the predefined situations, or you can create your own situations. Customizing predefined situations are not recommended, because the change can be overwritten by future software updates. Situations can involve more than one attribute. For example, you can create a situation that is true if the utilization of i-node for a file system is more than 95%, or the free space is less than 5%. When a situation is true, you can also configure a command to automate a response necessary for recovering from a certain unwanted condition.

Tivoli Enerprise Monitoring Server

The Tivoli Enerprise Monitoring Server is the key component on which all other architectural components depend directly. The Tivoli Enerprise Monitoring Server acts as a collection and control point for alerts received from the agents, and collects the performance and availability data from the agents.

The monitoring server is responsible for tracking the heartbeat request interval for all Tivoli Enterprise Monitoring Agents connected to it. The monitoring server stores, initiates, and tracks all situations and policies, and is the central repository for storing all active conditions on every Tivoli Enterprise Monitoring Agent. Additionally, it is responsible for initiating and tracking all generated actions that invoke a script or program on the Tivoli Enterprise Monitoring Agent.

The monitoring server storage repository is a proprietary database format, referred to as the Enterprise Information Base (EIB), grouped as a collection of files located on the Tivoli Enterprise Monitoring Server. The primary monitoring server is configured as a hub as shown in Figure 4-23 on page 196. All IBM Tivoli

Monitoring V6.2 installations require at least one monitoring server to be configured as a hub.

Additional remote monitoring servers can be used to create a scalable hub-spoke configuration into the architecture. This hub-remote interconnection provides a hierarchical design that enables the remote monitoring server to control and collect its individual agent status and propagate the agent status up to the hub monitoring server. This mechanism enables the hub monitoring server to maintain infrastructure-wide visibility of the entire environment.

Tivoli Enterprise Portal Server

The Tivoli Enterprise Portal Server (TEPS) is a repository for all graphical presentation of monitoring data. TEPS provides the core presentation layer, which allows for retrieval, manipulation, analysis, and reformatting of data. It manages this access through a portal client or Web browser.

Tivoli Enterprise Portal Client

The Tivoli Enterprise Portal Client, referred to as the portal client, is a Java™ user interface that connects to the Tivoli Enterprise Portal Server to display all monitoring data collections. You can run the portal client as a pre-installed standalone application, or as a Java Applet application. The portal client brings all of these views together in a single window so you can see when any component is not working as expected.

Tivoli Data Warehouse

With Tivoli Data Warehouse, you can analyze historical trends from monitoring agents. The Tivoli Data Warehouse uses a DB2, Oracle®, or Microsoft® SQL Server® database to store historical data collected across your environment. You can generate warehouse reports for short term or long term data through the Tivoli Enterprise Portal.

Warehouse Proxy agent

The Warehouse Proxy agent (WPA) is a unique agent that performs the task of receiving and consolidating all historical data collections from the individual agents to store in the Tivoli Data Warehouse.

Warehouse Summarization and Pruning agent

The Warehouse Summarization and Pruning (S&P) agent provides the ability to customize the length of time for which to save data (pruning) and how often to aggregate granular data (summarization).

For information about each component, see “2.1 IBM Tivoli Monitoring V6.2 components” in *Certification Study Guide Series: IBM Tivoli Monitoring V6.2*, SG24-7456

Tivoli Monitoring in Power Systems environment

As stated earlier, ITM covers the wide range of monitoring operations from a hardware level to an application level. The actual data collection functionality is implemented by monitoring agents. Various agents are designed specifically for Power Systems and their virtualization functions. This section describes the following monitoring agents:

- CEC Base Agent
- VIOS Premium Agent
- HMC Base Agent
- AIX Premium Agent

CEC Base Agent

Central Electronic Complex (CEC) Base Agent provides information of the availability and the resource utilization of Power Systems through CEC, which is the main hardware component of Power Systems hardware. This agent provides a global point of view for the availability and the resource utilization of Power Systems running multitenant workloads.

CEC Base Agent shows an inventory of CEC resources and resources allocated to individual LPARs on the CEC. The monitoring agent for CEC monitors the number of LPARs, CPU and memory allocations per LPAR, LPAR state, LPAR utilization, operating environment, CEC modes, and CEC utilization.

In this section, we show an example of a CEC resource inventory monitoring workspace. However, for details about the CEC Base Agent, see *IBM Tivoli Monitoring: CEC Base Agent User's Guide*, SC23-5239 (listed in the Tivoli Monitoring documentation tree as *Monitoring Agent for CEC Base User's Guide*):

http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.1/cec/om_net_agent_template.htm

In our example, the data is collected by the CEC Base Agent and is accessible by the portal client. In Figure 4-24 on page 200, the portal client shows the number of LPARs, CPU and memory allocation for these LPARs, and overall CPU and memory resource allocation status.

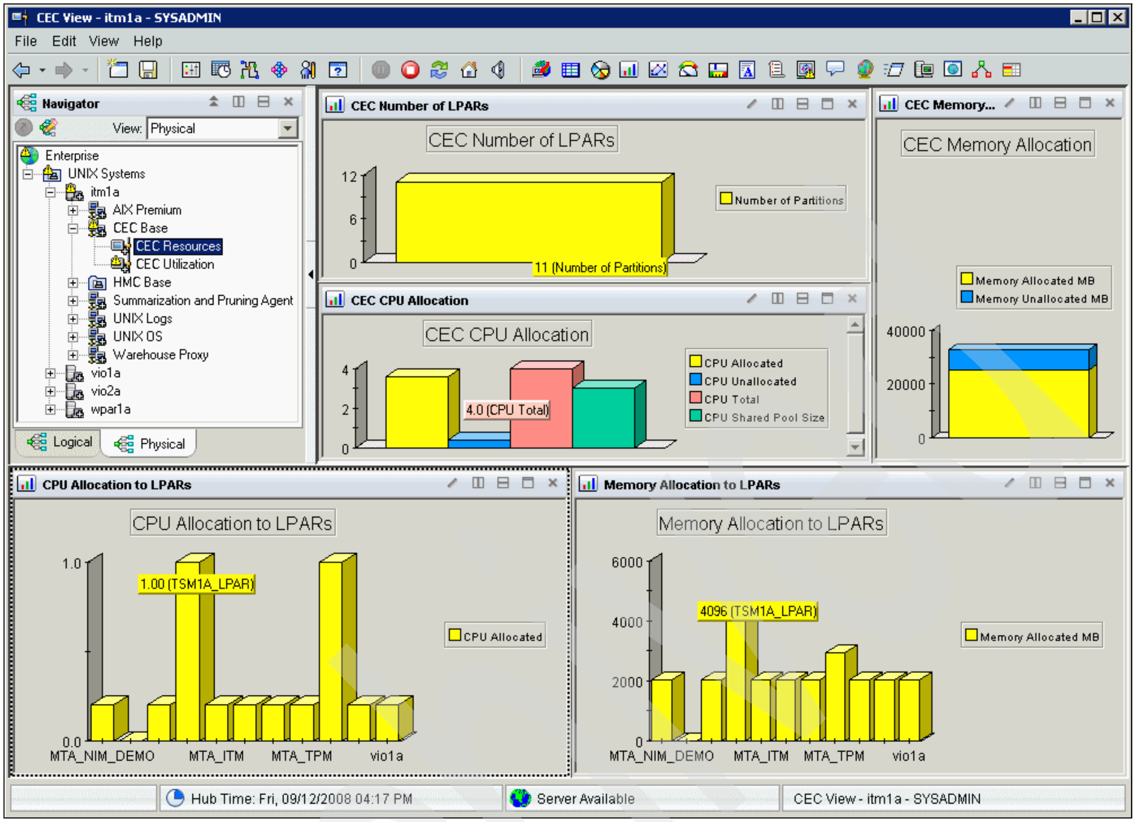

Figure 4-24 Base agent - overall resource allocation

The next window, shown in Figure 4-25 on page 201, has more detailed system resource allocation status for each LPAR. The *monitored systems* are the systems that are are running the xmtopas or xmservd daemon and are accessible from the CEC Base Agent.

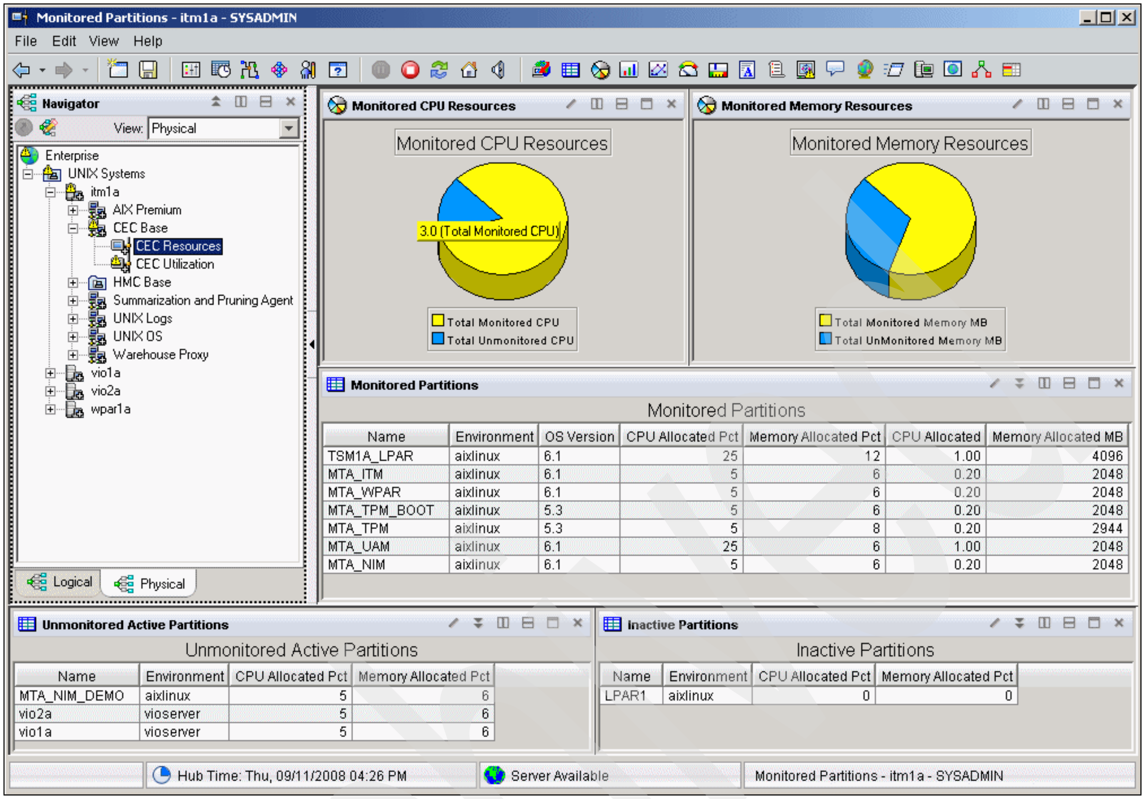

Monitored Partitions

Name	Environment	OS Version	CPU Allocated Pct	Memory Allocated Pct	CPU Allocated	Memory Allocated MB
TSM1A_LPAR	aixlinux	6.1	25	12	1.00	4096
MTA_ITM	aixlinux	6.1	5	6	0.20	2048
MTA_WPAR	aixlinux	6.1	5	6	0.20	2048
MTA_TPM_BOOT	aixlinux	5.3	5	6	0.20	2048
MTA_TPM	aixlinux	5.3	5	8	0.20	2944
MTA_UAM	aixlinux	6.1	25	6	1.00	2048
MTA_NIM	aixlinux	6.1	5	6	0.20	2048

Unmonitored Active Partitions

Name	Environment	CPU Allocated Pct	Memory Allocated Pct
MTA_NIM_DEMO	aixlinux	5	6
vio2a	vioserver	5	6
vio1a	vioserver	5	6

Inactive Partitions

Name	Environment	CPU Allocated Pct	Memory Allocated Pct
LPAR1	aixlinux	0	0

Figure 4-25 Base agent - Monitored Partitions view

Physical CPU and memory resource consumption by each LPAR is shown in Figure 4-26 on page 202.

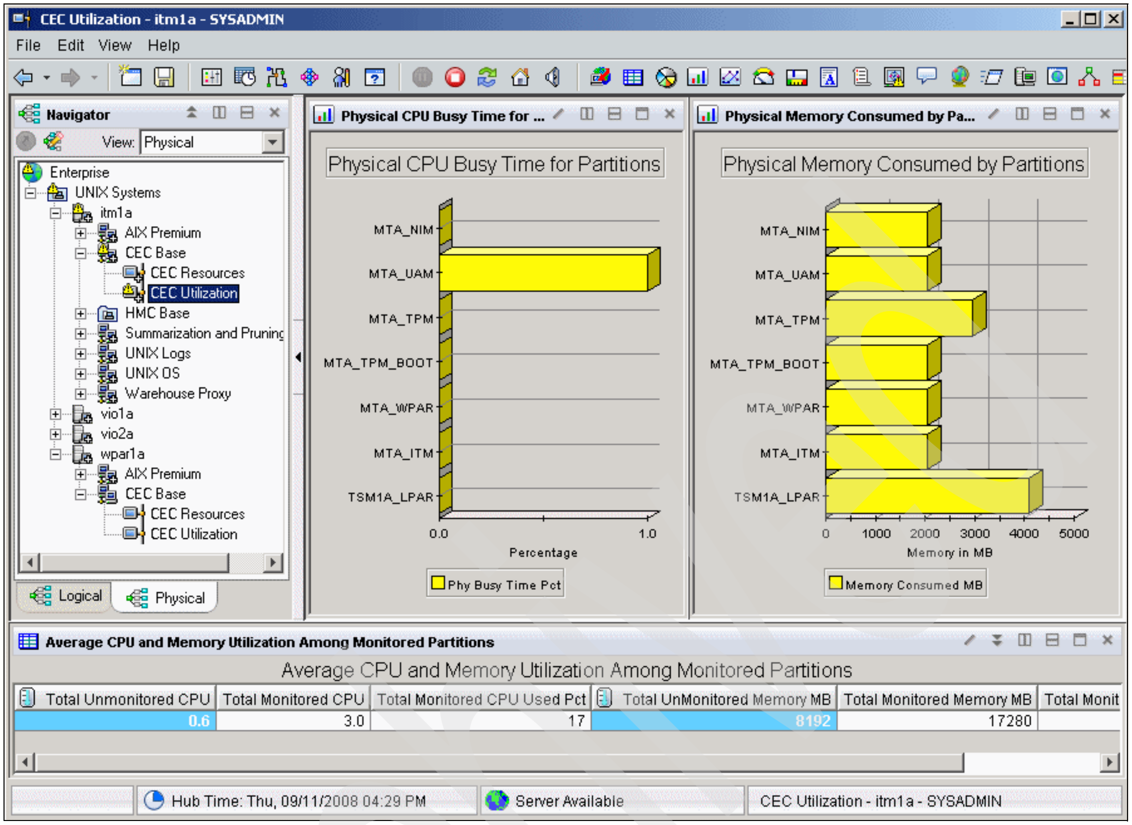

Average CPU and Memory Utilization Among Monitored Partitions

Total Unmonitored CPU	Total Monitored CPU	Total Monitored CPU Used Pct	Total UnMonitored Memory MB	Total Monitored Memory MB	Total Monit
0.6	3.0	17	8192	17280	

Figure 4-26 Base agent - physical CPU and memory resource consumption

When you use CEC Base Agent, you have to install the agent to a server that is running supported operating systems, which can be AIX 5.3 TL5 or later, or AIX 6.1 system. The server must be able to have the SSH connection to the Hardware Management Console, which manages the Power Systems servers you want. This agent should not be installed onto the tenant system, because CEC Base Agent has the global view of the whole system.

> **Note:** A CEC Base Agent can monitor a single managed Power Systems server. To monitor multiple Power Systems, you have to install additional CEC agents in separate operating systems managed by the service provider.

CEC Base Agent can collect resource utilization data such as the CPU utilization of tenant LPARs. For this, you should have xmtopas or xmservd daemon running on the tenant LPARs. You also have to configure the network to allow the connections, as listed in Table 4-4 on page 203.

Table 4-4 Required network connection for CEC agent

Service	Source	Destination	SRC port	DST port
SSH	CEC agent	HMC	Any	22:tcp
xmquery	CEC agent	tenant LPAR	Any	2279:tcp 2279:udp

In addition to the ports listed in the table, CEC Base Agent uses a network connection to the monitoring server as described in “Tivoli Monitoring Server and Agent firewall considerations” on page 211.

VIOS Premium Agent

The VIOS Premium Agent monitors the availability and health of the Virtual I/O Server resources. Virtual I/O Server, one of PowerVM components, contains the VIOS Base Agent by default, and when you install VIOS onto an LPAR, the VIOS Base Agent is installed automatically. The VIOS Premium Agent is an optional monitoring component for VIOS and can collect much more information about the VIOS activity at high granularity.

For more information, see *IBM Tivoli Monitoring: VIOS Premium Agent User’s Guide*, SC23-2238 (listed in the Tivoli Monitoring documentation tree as *Monitoring Agent for VIOS Premium User's Guide*):

http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.1/VIOS_Premium_Agent.htm

VIOS Premium Agent monitors CPU, memory, storage and networks performance measurements of the Virtual I/O Server. It also shows storage and network mappings between the Virtual I/O Server and its clients.

Figure 4-27 on page 204 shows the historical CPU usage, the memory consumption, the network, and the I/O transfer ratio for Virtual I/O Server. To obtain the historical data, you also have to set up the Tivoli Data Warehouse and the Warehouse Proxy agent.

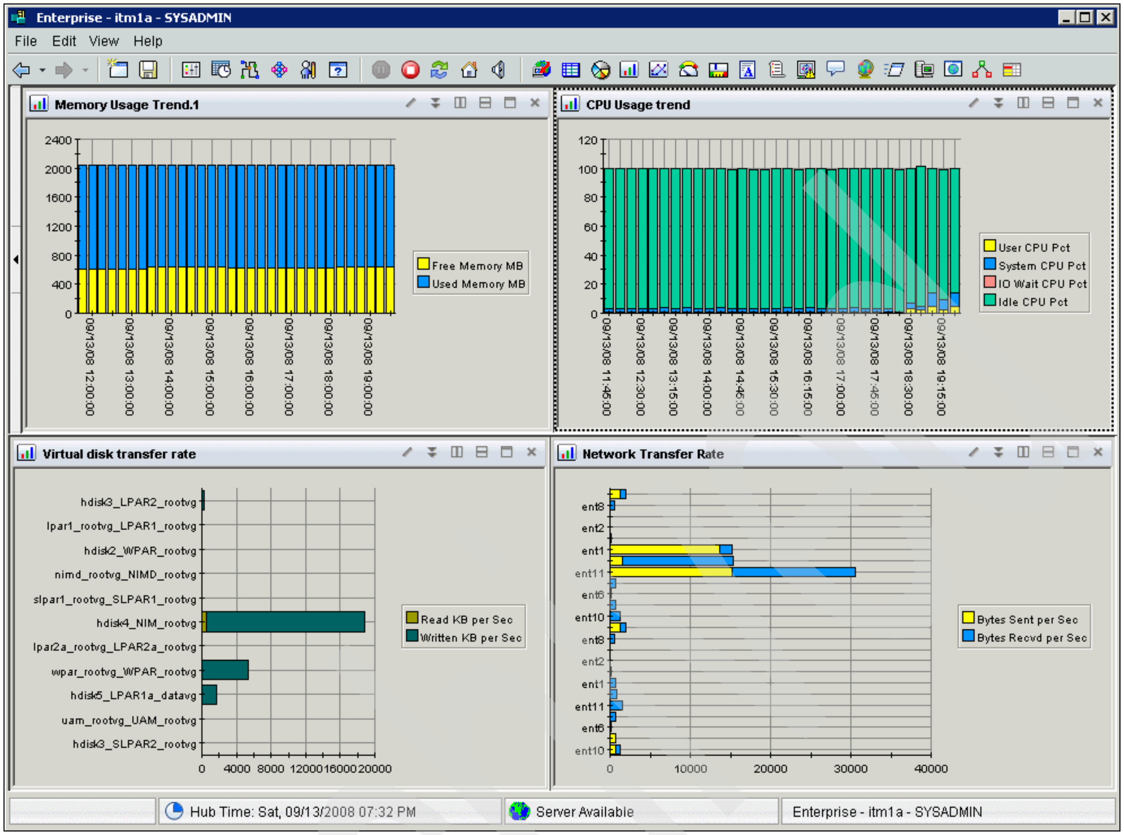

Figure 4-27 Virtual I/O Server performance data

Figure 4-28 on page 205 shows storage mapping information about Virtual I/O Server with VIOS Premium Agent.

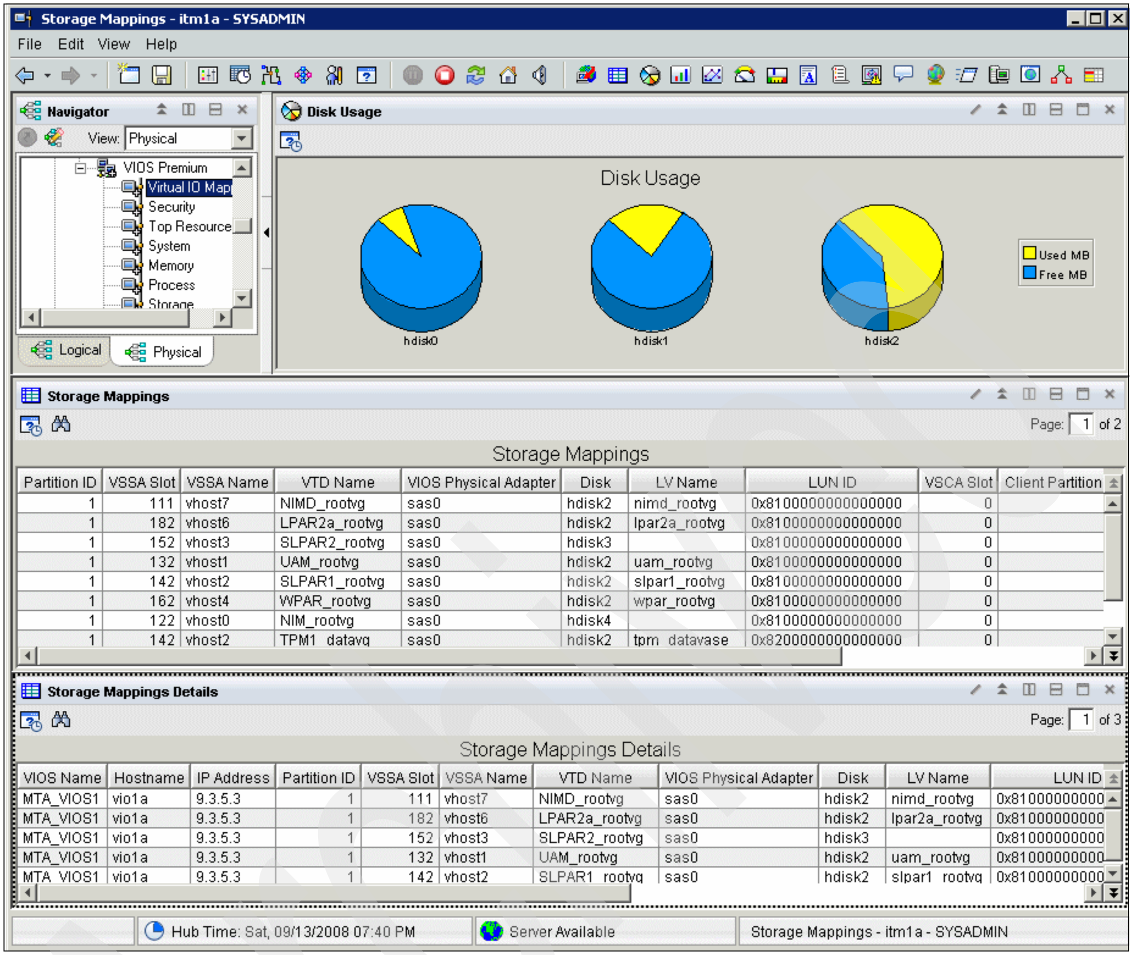

Storage Mappings

Partition ID	VSSA Slot	VSSA Name	VTD Name	VIOS Physical Adapter	Disk	LV Name	LUN ID	VSCA Slot	Client Partition
1	111	vhost7	NIMD_rootvg	sas0	hdisk2	nimd_rootvg	0x8100000000000000	0	
1	182	vhost6	LPAR2a_rootvg	sas0	hdisk2	lpar2a_rootvg	0x8100000000000000	0	
1	152	vhost3	SLPAR2_rootvg	sas0	hdisk3		0x8100000000000000	0	
1	132	vhost1	UAM_rootvg	sas0	hdisk2	uam_rootvg	0x8100000000000000	0	
1	142	vhost2	SLPAR1_rootvg	sas0	hdisk2	slpar1_rootvg	0x8100000000000000	0	
1	162	vhost4	WPAR_rootvg	sas0	hdisk2	wpar_rootvg	0x8100000000000000	0	
1	122	vhost0	NIM_rootvg	sas0	hdisk4		0x8100000000000000	0	
1	142	vhost2	TPM1_datavg	sas0	hdisk2	tpm_datavase	0x8200000000000000	0	

Storage Mappings Details

VIOS Name	Hostname	IP Address	Partition ID	VSSA Slot	VSSA Name	VTD Name	VIOS Physical Adapter	Disk	LV Name	LUN ID
MTA_VIOS1	vio1a	9.3.5.3	1	111	vhost7	NIMD_rootvg	sas0	hdisk2	nimd_rootvg	0x81000000000
MTA_VIOS1	vio1a	9.3.5.3	1	182	vhost6	LPAR2a_rootvg	sas0	hdisk2	lpar2a_rootvg	0x81000000000
MTA_VIOS1	vio1a	9.3.5.3	1	152	vhost3	SLPAR2_rootvg	sas0	hdisk3		0x81000000000
MTA_VIOS1	vio1a	9.3.5.3	1	132	vhost1	UAM_rootvg	sas0	hdisk2	uam_rootvg	0x81000000000
MTA_VIOS1	vio1a	9.3.5.3	1	142	vhost2	SLPAR1_rootvg	sas0	hdisk2	slpar1_rootvg	0x81000000000

Figure 4-28 Virtual I/O ServerStorage Mapping information

HMC Base Agent

HMC Base Agent monitors the availability and health of the Hardware Management Console (HMC) resources such as CPU, memory, storage, and network. HMC Base Agent also shows the information about the HMC configuration, the managed systems, and the LPARs. Figure 4-29 on page 206 shows an example of the data collected by the HMC Base Agent.

For more information, see *IBM Tivoli Monitoring: HMC Base Agent User's Guide:*

http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.1/HMC_Base_Agent.htm

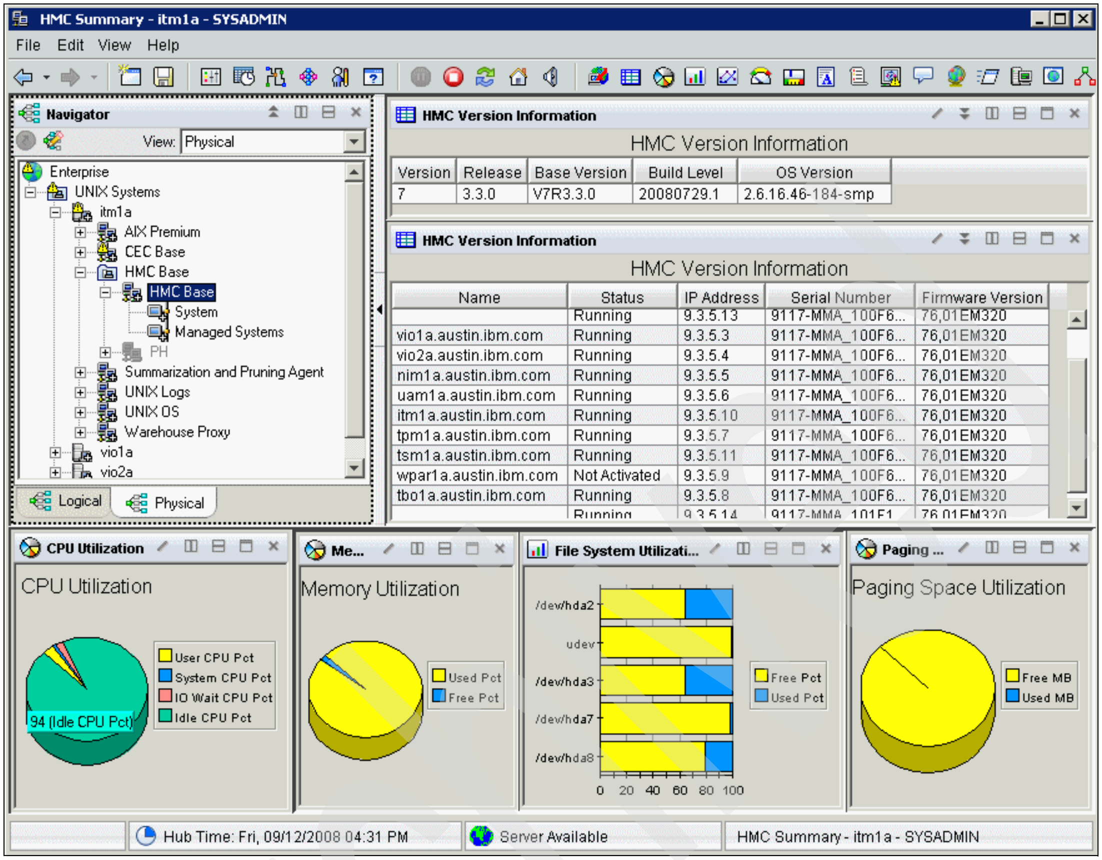

HMC Version Information

Version	Release	Base Version	Build Level	OS Version
7	3.3.0	V7R3.3.0	20080729.1	2.6.16.46-184-smp

HMC Version Information

Name	Status	IP Address	Serial Number	Firmware Version
	Running	9.3.5.13	9117-MMA_100F6...	76,01EM320
vio1a.austin.ibm.com	Running	9.3.5.3	9117-MMA_100F6...	76,01EM320
vio2a.austin.ibm.com	Running	9.3.5.4	9117-MMA_100F6...	76,01EM320
nim1a.austin.ibm.com	Running	9.3.5.5	9117-MMA_100F6...	76,01EM320
uam1a.austin.ibm.com	Running	9.3.5.6	9117-MMA_100F6...	76,01EM320
itm1a.austin.ibm.com	Running	9.3.5.10	9117-MMA_100F6...	76,01EM320
tpm1a.austin.ibm.com	Running	9.3.5.7	9117-MMA_100F6...	76,01EM320
tsm1a.austin.ibm.com	Running	9.3.5.11	9117-MMA_100F6...	76,01EM320
wpar1a.austin.ibm.com	Not Activated	9.3.5.9	9117-MMA_100F6...	76,01EM320
tbo1a.austin.ibm.com	Running	9.3.5.8	9117-MMA_100F6...	76,01EM320
	Running	9.3.5.14	9117-MMA_101F1	76,01EM320

Figure 4-29 HMC Base Agent - HMC performance and managed systems information

AIX Premium Agent

AIX Premium Agent can monitor the availability, health, and performance of key AIX system resources such as LPAR configurations, CPU, memory, storage, network, printers, NIM, and WPARs performance metrics.

We show an example of information monitored by AIX Premium Agent in this section. However, for more details, see the *IBM Tivoli Monitoring: AIX Premium Agent User's Guide*, SC23-2237 (listed in the Tivoli Monitoring documentation tree as *Monitoring Agent for AIX Premium User's Guide*):

http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc_6.1/aix/om_net_agent_template.htm

Figure 4-30 on page 207 shows an example of the CPU utilization of an AIX system.

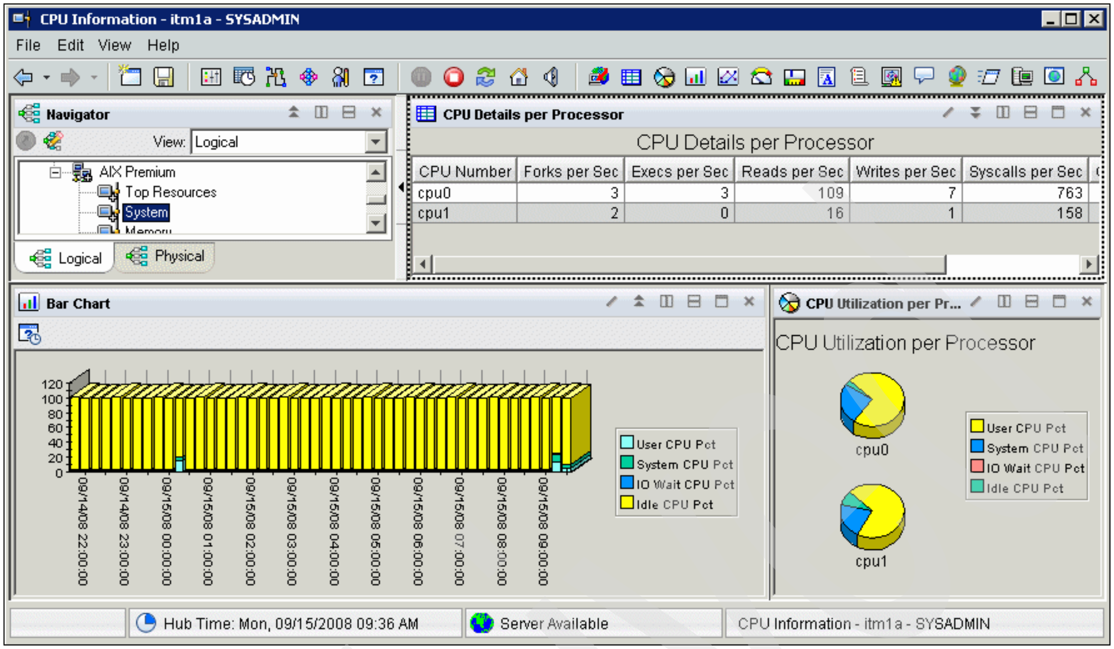

CPU Details per Processor

CPU Number	Forks per Sec	Execs per Sec	Reads per Sec	Writes per Sec	Syscalls per Sec
cpu0	3	3	109	7	763
cpu1	2	0	16	1	158

Figure 4-30 AIX Premium Agent - CPU utilization information

Figure 4-31 on page 208 shows the memory consumption information of an AIX system.

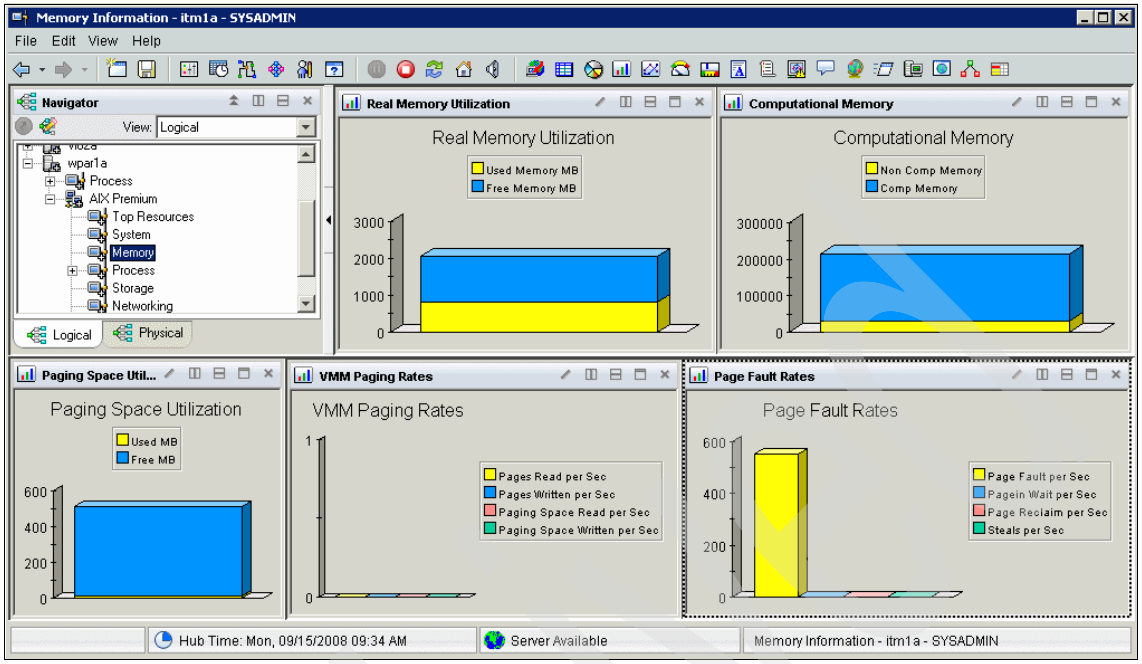

Figure 4-31 AIX Premium Agent - memory consumption information

Figure 4-32 on page 209 shows the network interface utilization.

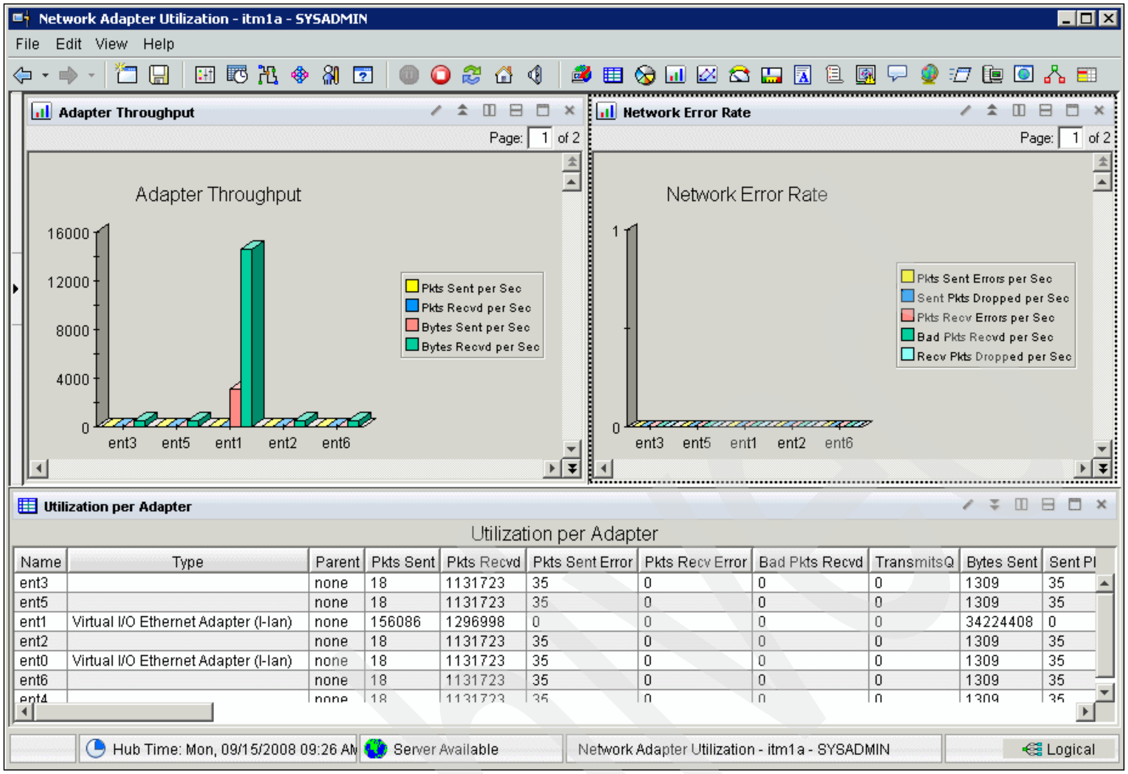

Utilization per Adapter

Name	Type	Parent	Pkts Sent	Pkts Recvd	Pkts Sent Error	Pkts Recv Error	Bad Pkts Recvd	TransmitsQ	Bytes Sent	Sent PI
ent3		none	18	1131723	35	0	0	0	1309	35
ent5		none	18	1131723	35	0	0	0	1309	35
ent1	Virtual I/O Ethernet Adapter (l-lan)	none	156086	1296998	0	0	0	0	34224408	0
ent2		none	18	1131723	35	0	0	0	1309	35
ent0	Virtual I/O Ethernet Adapter (l-lan)	none	18	1131723	35	0	0	0	1309	35
ent6		none	18	1131723	35	0	0	0	1309	35
ent4		none	18	1131723	35	0	0	0	1309	35

Hub Time: Mon, 09/15/2008 09:26 AM
Server Available
Network Adapter Utilization - itm1a - SYSADMIN
Logical

Figure 4-32 AIX Premium Agent - network utilization information

Figure 4-33 on page 210 shows the file system utilization, and top processes information.

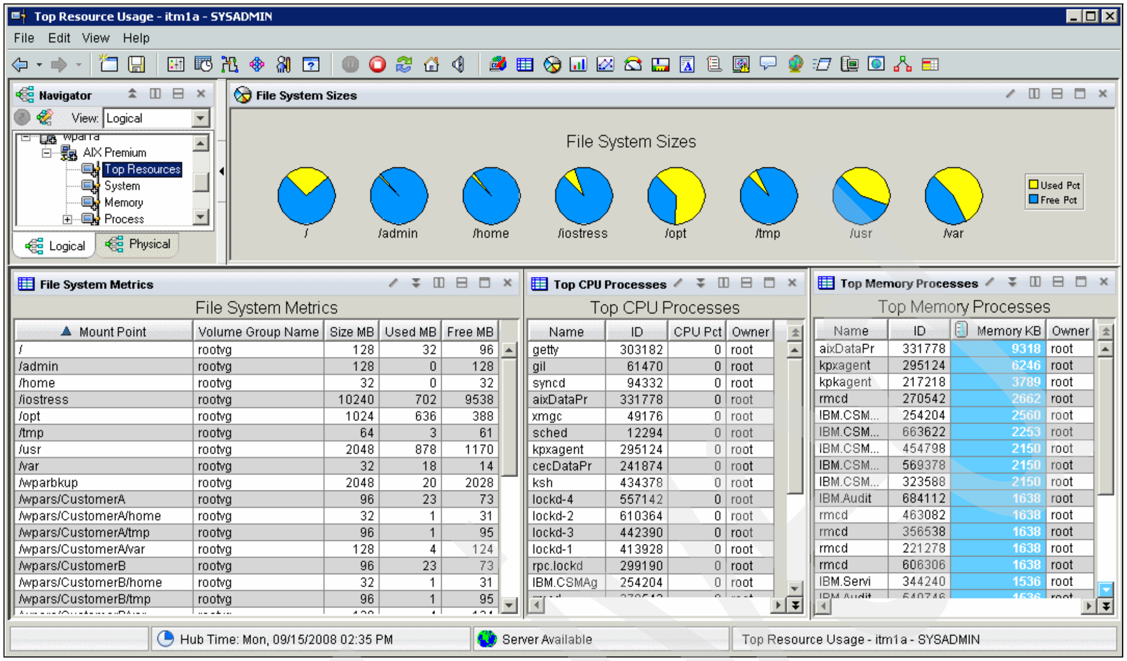

File System Metrics

Mount Point	Volume Group Name	Size MB	Used MB	Free MB
/	rootvg	128	32	96
/admin	rootvg	128	0	128
/home	rootvg	32	0	32
/iostress	rootvg	10240	702	9538
/opt	rootvg	1024	636	388
/tmp	rootvg	64	3	61
/usr	rootvg	2048	878	1170
/var	rootvg	32	18	14
/wparbkup	rootvg	2048	20	2028
/wpars/CustomerA	rootvg	96	23	73
/wpars/CustomerA/home	rootvg	32	1	31
/wpars/CustomerA/tmp	rootvg	96	1	95
/wpars/CustomerA/var	rootvg	128	4	124
/wpars/CustomerB	rootvg	96	23	73
/wpars/CustomerB/home	rootvg	32	1	31
/wpars/CustomerB/tmp	rootvg	96	1	95

Top CPU Processes

Name	ID	CPU Pct	Owner
getty	303182	0	root
gil	61470	0	root
syncd	94332	0	root
aixDataPr	331778	0	root
xmgc	49176	0	root
sched	12294	0	root
kpxagent	295124	0	root
cecDataPr	241874	0	root
ksh	434378	0	root
lockd-4	557142	0	root
lockd-2	610364	0	root
lockd-3	442390	0	root
lockd-1	413928	0	root
rpc.lockd	299190	0	root
IBM.CSMAg	254204	0	root

Top Memory Processes

Name	ID	Memory KB	Owner
aixDataPr	331778	9318	root
kpxagent	295124	6246	root
kpkagent	217218	3789	root
rmcd	270542	2662	root
IBM.CSM...	254204	2560	root
IBM.CSM...	663622	2253	root
IBM.CSM...	454798	2150	root
IBM.CSM...	569378	2150	root
IBM.CSM...	323588	2150	root
IBM.Audit	684112	1638	root
rmcd	463082	1638	root
rmcd	356538	1638	root
rmcd	221278	1638	root
rmcd	606306	1638	root
IBM.Servi	344240	1536	root

Hub Time: Mon, 09/15/2008 02:35 PM

Server Available

Top Resource Usage - itm1a - SYSADMIN

Figure 4-33 AIX Premium Agent - file system utilization and top process information

Figure 4-34 on page 211 shows the WPAR utilization information of an AIX system.

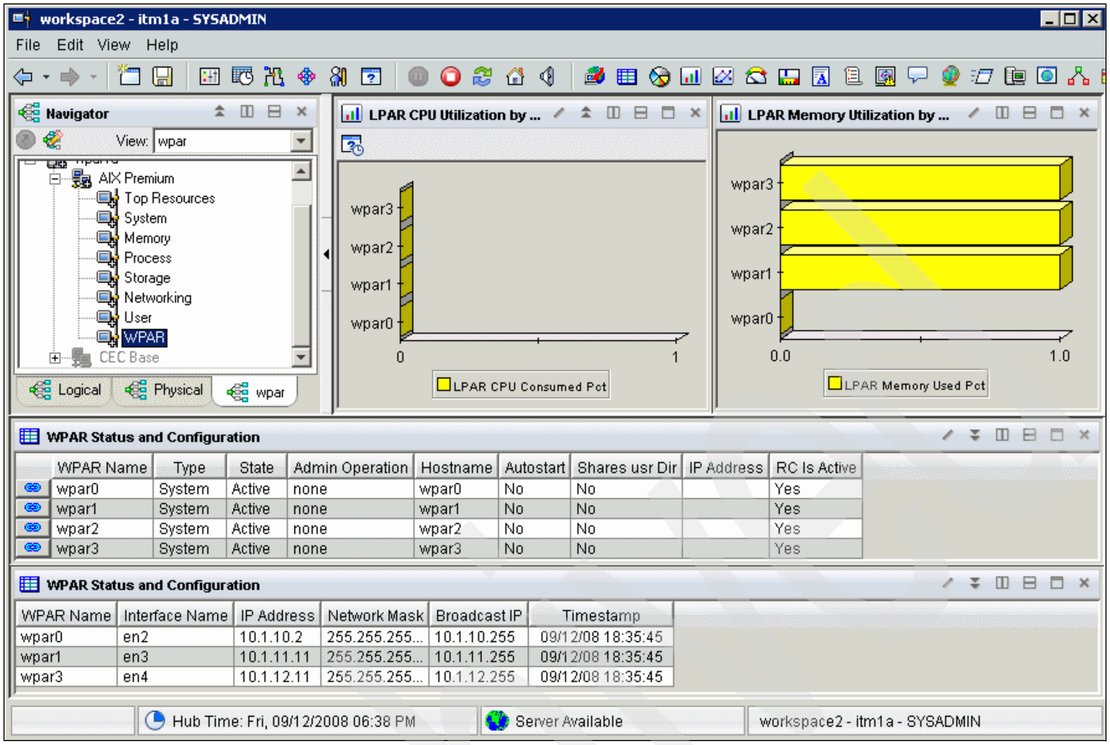

WPAR Status and Configuration

WPAR Name	Type	State	Admin Operation	Hostname	Autostart	Shares usr Dir	IP Address	RC Is Active
wpar0	System	Active	none	wpar0	No	No		Yes
wpar1	System	Active	none	wpar1	No	No		Yes
wpar2	System	Active	none	wpar2	No	No		Yes
wpar3	System	Active	none	wpar3	No	No		Yes

WPAR Status and Configuration

WPAR Name	Interface Name	IP Address	Network Mask	Broadcast IP	Timestamp
wpar0	en2	10.1.10.2	255.255.255...	10.1.10.255	09/12/08 18:35:45
wpar1	en3	10.1.11.11	255.255.255...	10.1.11.255	09/12/08 18:35:45
wpar3	en4	10.1.12.11	255.255.255...	10.1.12.255	09/12/08 18:35:45

Figure 4-34 AIX Premium Agent - WPAR information

Tivoli Monitoring Server and Agent firewall considerations

In a multitenant environment, the network security must be considered. Especially critical to consider is the network connection between the monitoring server and the monitoring agents which are installed on tenant system. When you use IBM Tivoli Monitoring (ITM) on a Power Systems environment, three network services are available for connection between the monitoring server and the agents. They are IP(UDP), IP.PIPE(TCP), or IP.SPIPE(TCP with SSL). If you are installing ITM components across firewalls, configure each component with the IP.PIPE or IP.SPIPE services. Table 4-5 lists ports used for the connection.

Table 4-5 Required ports between the monitoring agents and the monitoring server

Service	Source	Destination	SRC port	DST port
IP.PIPE	Monitoring agents	Monitoring server	Any	1918:tcp
IP.SPIPE (SSL)	Monitoring agents	Monitoring server	Any	3660:tcp

For further description of the firewall between monitoring servers and agents, see “3.1 IBM Tivoli Monitoring network components and flows” in *IBM Tivoli Monitoring: Implementation and Performance Optimization for Large Scale Environments*, SG24-7443.

TEPS and portal client firewall considerations

The connection between Tivoli Enterprise Portal Server (TEPS) and the portal clients are HTTP, although the portal server can support SSL connections, too. To configure the portal server for the SSL capability, see the following resources:

- Additional Tivoli Enterprise Portal Configuration

  http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/topic/com.ibm.itm.doc_6.2.fp1/itm_install184.htm

- Firewall scenarios for Tivoli Enterprise Portal

  http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc_6.2.fp1/itm_install199.htm

Tivoli Monitoring requirement for monitored systems

With few exceptions, all monitored systems require that the Tivoli Enerprise Monitoring Agent be installed so that they can be monitored by Tivoli Enerprise Monitoring Server. You have to install, configure, and start the agent.

Customizing Tivoli Enterprise Portal

Tivoli Enterprise Portal is highly customizable with simple GUI operations. You can change the type of chart, convert it to a table, choose the items shown in the chart, set a filter for the result, or create your own queries for desired monitoring data. For detailed procedures of customizing the views, see: “7.2 Workspace” in *Certification Study Guide Series: IBM Tivoli Monitoring V6.2*, SG24-7456, and “Chapter 6, Customizing workspace” in *Tivoli Monitoring User’s Guide,* SC32-9409.

Historical information access

Tivoli Monitoring stores the historical monitoring data in Tivoli Data Warehouse. The historical data is accessible and can be displayed through the portal client, or any other reporting tool such as Business Intelligence and Reporting Tools (BIRT), which has JDBC™ connectivity. For further information about historical data in Tivoli Monitoring,see “Chapter 4, Planning historical data collection in large scale environments” and “Chapter 11, Reporting enhancements” in *IBM Tivoli Monitoring: Implementation and Performance Optimization for Large Scale Environments*, SG24-7443.

IBM Tivoli Monitoring publications

The following resources can help you understand, install, and configure Tivoli Monitoring products:

- For a general overview, see *IBMTivoli Monitoring*, GC28-8379:

  ftp://ftp.software.ibm.com/software/tivoli/datasheets/TivoliMonitoring61_GC28-8379-01.pdf

- The IBM Tivoli Monitoring Web site:

  http://www.ibm.com/software/tivoli/products/monitor/

IBM Tivoli Monitoring publications are as follows:

- *Certification Study Guide Series: IBM Tivoli Monitoring V6.2*, SG24-7456
- *IBM Tivoli Monitoring: Implementation and Performance Optimization for Large Scale Environments*, SG24-7443
- *Deployment Guide Series: IBM Tivoli Monitoring V6.2*, SG24-7444

IBM Tivoli Monitoring manuals are:

- *IBM Tivoli Monitoring: Installation and Setup Guide*, GC32-9407
- *IBM Tivoli Monitoring: Administrator's Guide*, SC32-9408
- *IBM Tivoli Monitoring: User's Guide*, SC32-9409

The IBM Tivoli Monitoring documentation is located on the following Web site:

http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/index.jsp?topic=/com.ibm.itm.doc/welcome.htm

4.4.4 IBM Director

IBM Director provides an integrated suite of software tools for a consistent, single point of management and automation. IBM Director provides clients with flexible capabilities to realize maximum system availability and lower IT costs. With IBM Director, IT administrators can view and track the hardware configuration of remote systems in detail and monitor the usage and performance of critical components, such as processors, disks, and memory.

IBM Director enables monitoring and event management across a heterogeneous IT environment, including AIX, IBM i, Windows®, Linux, NetWare, and VMware® ESX Server, from a single user interface. From one access point, administrators can monitor system resources, inventory, and events. They can also perform task management, core corrective actions, distributed commands, and hardware control for servers, clients, and storage. We discuss IBM Director V5.2, unless stated otherwise.

IBM Director components

IBM Director is designed to manage a complex environment that contains numerous servers, client computers, storage subsystems, and SNMP-based devices. IBM Director consists of various components, and the combination of component vary for monitoring purposes. Figure 4-35 shows an example topology of IBM Director components for monitoring and managing Power Systems servers. LPARs configured in Power Systems can be managed with or without client modules installed. We describe each IBM Director component.

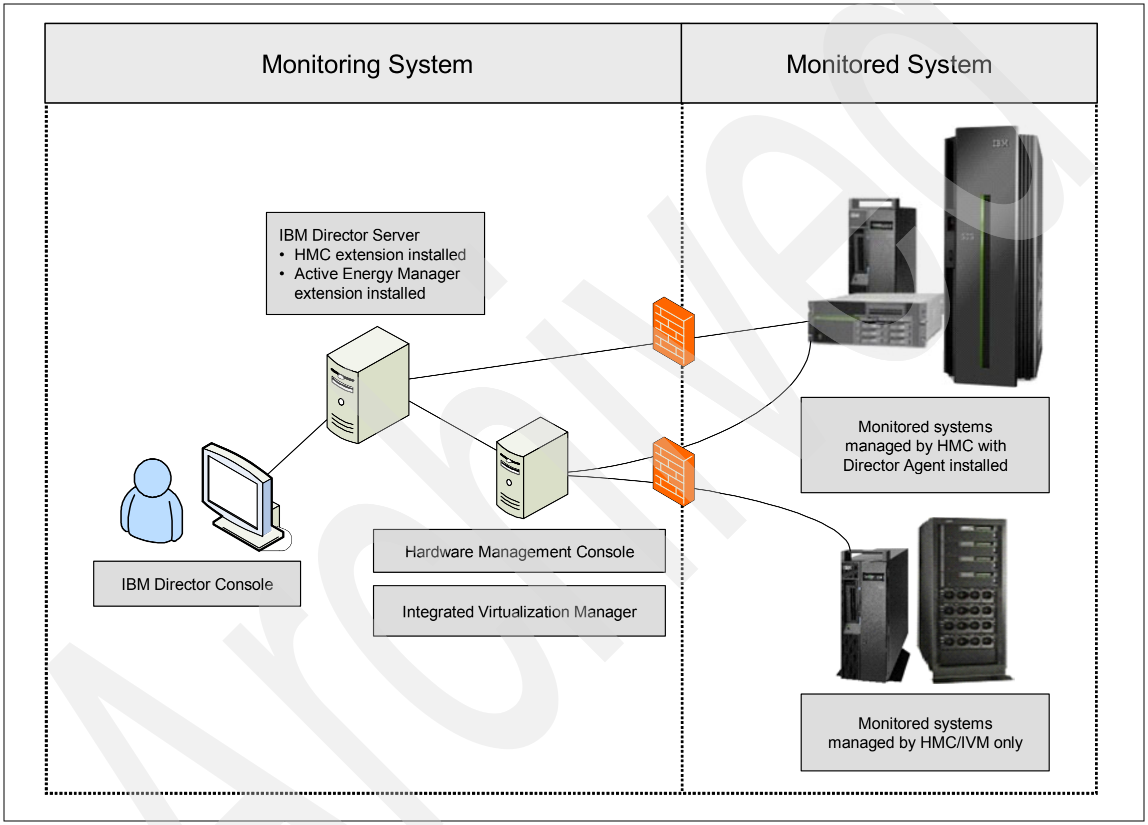
Monitoring System
Monitored System
IBM Director Server
• HMC extension installed
• Active Energy Manager extension installed
IBM Director Console
Hardware Management Console
Integrated Virtualization Manager
Monitored systems managed by HMC with Director Agent installed
Monitored systems managed by HMC/IVM only

Figure 4-35 Monitoring and managing Power Systems with IBM Director

IBM Director Server

IBM Director Server is the main component of IBM Director. IBM Director Server contains the management data, the server engine, and the application logic. It provides basic functions such as discovery of managed systems, persistent storage of inventory data, relational database support, presence-checking, security and authentication, management console support, and administrative tasks.

IBM Director Console

IBM Director Console is the graphical user interface (GUI) for IBM Director Server, as shown in Figure 4-36. With IBM Director Console, system administrators can conduct comprehensive hardware management by using either a drag-and-drop action or a single click.

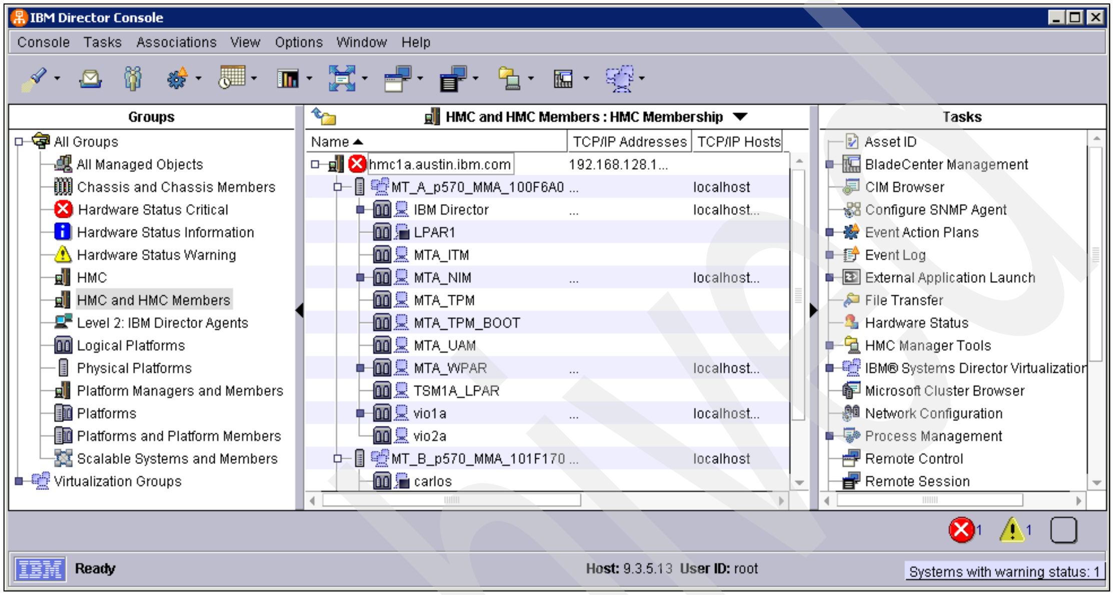

Name	TCP/IP Addresses	TCP/IP Hosts
hmc1a.austin.ibm.com	192.168.128.1...	
MT_A_p570_MMA_100F6A0	...	localhost
IBM Director	...	localhost...
LPAR1		
MTA_ITM		
MTA_NIM	...	localhost...
MTA_TPM		
MTA_TPM_BOOT		
MTA_UAM		
MTA_WPAR	...	localhost...
TSM1A_LPAR		
vio1a	...	localhost...
vio2a		
MT_B_p570_MMA_101F170	...	localhost
carlos		

Host: 9.3.5.13 User ID: root

Systems with warning status: 1

Figure 4-36 Example of IBM Director Console managing Power Systems

IBM Director Agent

IBM Director Agent is installed on a managed system to provide enhanced functionality for IBM Director to communicate with and administer the managed system. IBM Director Agent provides management data to the management server through various network protocols.

IBM Director extensions

IBM Director can extend its capability with additional software called extension. For the purpose of monitoring Power Systems servers, the following extensions can be useful:

- Active Energy Manager extension

  Active Energy Manager is an extension of IBM Systems Director and is available for installation on Linux for Power System servers, Linux on x86 architecture, Linux on System z®, and Microsoft Windows. Active Energy Manager helps you monitor and manage the power usage of systems across IBM systems and other systems. Active Energy Manger is an energy management software tool that can provide a single view of the actual power usage across multiple platforms. It can effectively monitor and control power in the data center at the system, chassis, or rack level. By enabling these power management technologies, system administrators can more effectively power manage their systems while lowering the cost of computing.

Active Energy Manager extension provides the following functions:

– Power trending and thermal trending

  Power trending and thermal trending allow you to monitor the power consumption and temperature of a supported managed objects. Figure 4-37 on page 217 shows an example of the power and temperature trending monitoring capability.

> **Note:** Figure 4-37 is a sample taken from the beta version of Active Energy Manager V4.1 and IBM Director v6.1. The window might appear different in general availability versions because this section was based on an internal release of the code.

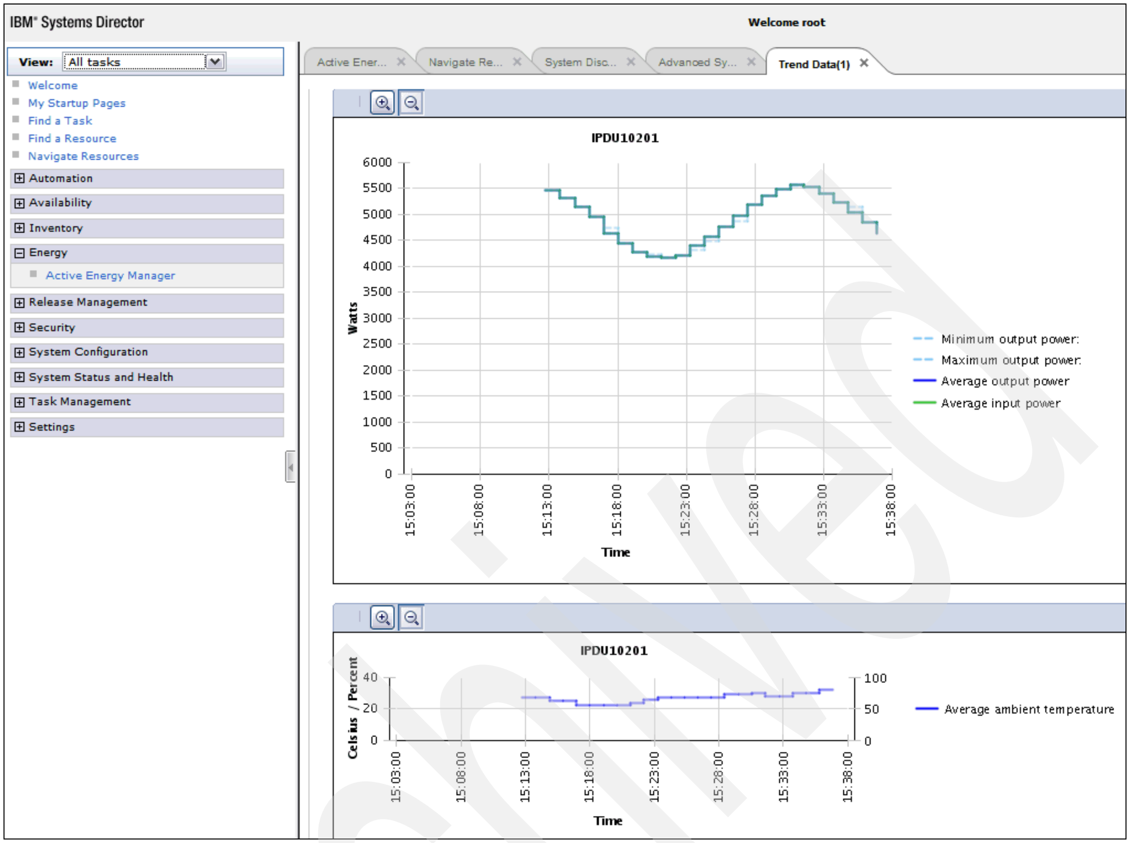

Figure 4-37 Example of power consumption and temperature trending

- CPU trending

  CPU trending helps you determine the actual CPU speed of processors for either the active power saver or active power cap function.

- Power saver

  Power saver helps you save energy when you throttle back the processor voltage and clocking rate. Use the power saver function to match computing power to workload while reducing your energy costs.

- Power capping

  Power capping helps you allocate less energy for a system when you set a cap on the number of watts that the power managed system can consume. If the power consumption of the server approaches the cap, Active Energy Manager throttles back the processor voltage and clocking rate in the same way as for the power saver function. In this way you can guarantee that the power cap value is not exceeded.

To enable and use the power management functionality, your systems has to be supported by Active Energy Manager. Active Energy Manager can monitor various servers such System x®, Power Systems, BladeCenter, System z, and intelligent power distribution units. For supported systems information, see the Managed Systems (for Active Energy Manager) on the IBM Systems Software Information Center:

http://publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/aem_310/frb0_r_HW_reqs_managed_systems.html

For details, see *Going Green with IBM Systems Director Active Energy Manager*, REDP-4361.

- Hardware Management Console extension

  Hardware Management Console (HMC) extension integrates the management and monitoring capability which usually HMC and IVM provides, into IBM Director framework. With the HMC extension installed, you can use the IBM Director console as a single point of management. HMC extension can perform the following functions:

  – Discover HMC, IVM, and Flexible Service Processor (FSP)-managed physical platforms by using industry-standard Service Location Protocol (SLP).

  – Manage HMC, IVM, and FSP-managed physical platforms with Common Information Model (CIM).

  – Show relationships among managed HMCs, IVMs, Power Systems servers, LPARS, and installed operating systems.

  – Perform power management tasks on all related managed platforms.

  – Collect hardware inventory for HMC and IVM-managed systems.

  – Receive hardware status and alerts in IBM Director Console.

  – Launch HMC console from IBM Director Console to perform advanced management tasks.

  The HMC extension is a prerequisite for the Active Energy Manager extension, described earlier. For information about HMC resources, see:

  http://publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/hmc_5.20/frk0_main.html

IBM Director monitoring features

This section provides examples of monitoring functionality of IBM Director in a Power Systems environment.

System availability monitoring

IBM Director offers system Power Systems failure detection feature provided by the Hardware Management Console extension. IBM Director integrates the system failure events, called serviceable events, to IBM Director framework and shows the IBM Director Console window. In Figure 4-36 on page 215, the symbol next to the HMC icon indicates that the systems managed by the HMC have serviceable events. To show the list of events, right-click the HMC icon and select **Hardware Status**, as shown in Figure 4-38.

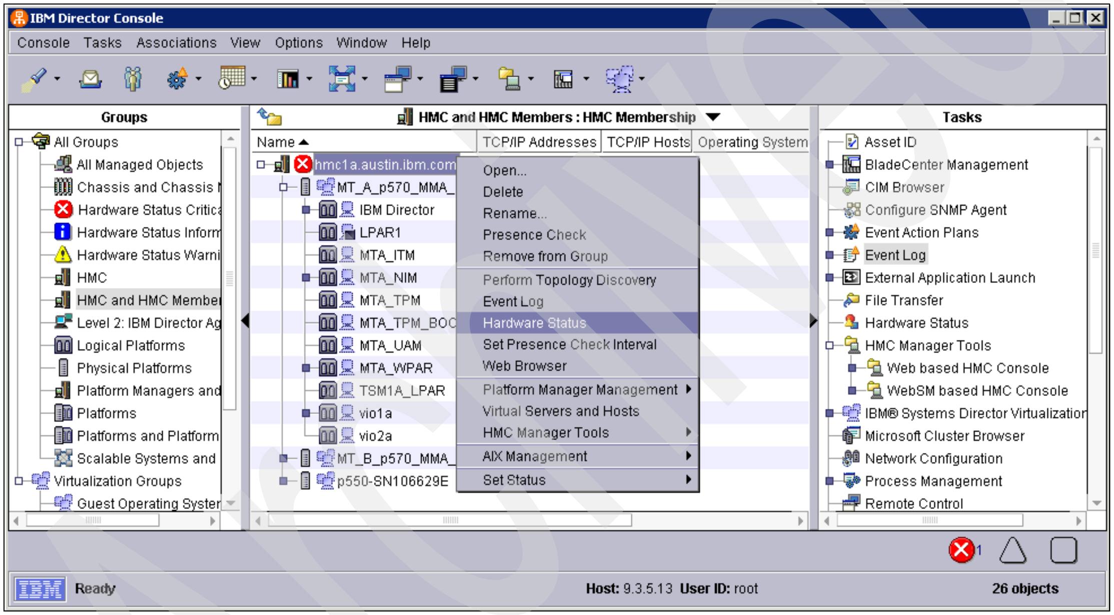

Figure 4-38 How to show the list of serviceable events

Figure 4-39 on page 220 provides an example of serviceable event list.

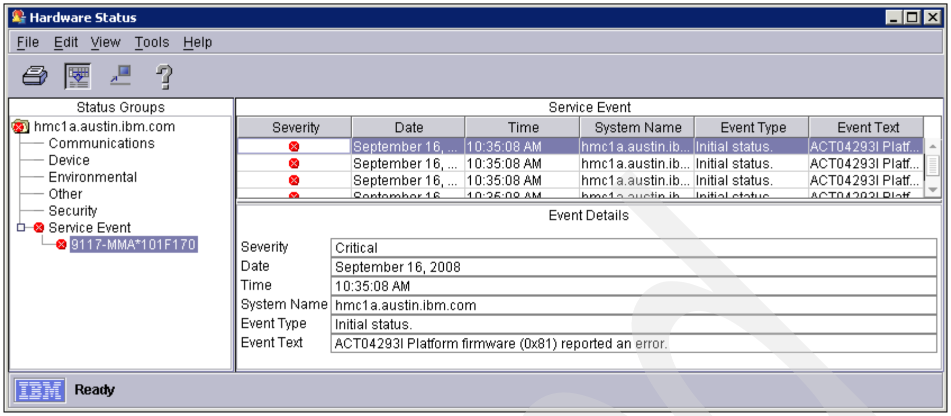

Service Event

Severity	Date	Time	System Name	Event Type	Event Text
⊗	September 16, ...	10:35:08 AM	hmc1a.austin.ib...	Initial status.	ACT04293I Platf...
⊗	September 16, ...	10:35:08 AM	hmc1a.austin.ib...	Initial status.	ACT04293I Platf...
⊗	September 16, ...	10:35:08 AM	hmc1a.austin.ib...	Initial status.	ACT04293I Platf...

Event Details

Severity	Critical
Date	September 16, 2008
Time	10:35:08 AM
System Name	hmc1a.austin.ibm.com
Event Type	Initial status.
Event Text	ACT04293I Platform firmware (0x81) reported an error.

Figure 4-39 HMC serviceable events shown by IBM Director

System resource utilization monitoring

IBM Director can monitor Power Systems server resource utilization measurements, such as CPU usage, memory consumption, network read-to-write ratio, and percentage of free space in file systems. Monitoring these objects requires the monitored systems to be the IBM Director Agent software installed. IBM Director can show real-time resource utilization, or the historical view.

To start the system resource monitoring window, right-click on the monitored system and select **Resource Monitors**, as shown in Figure 4-40 on page 221.

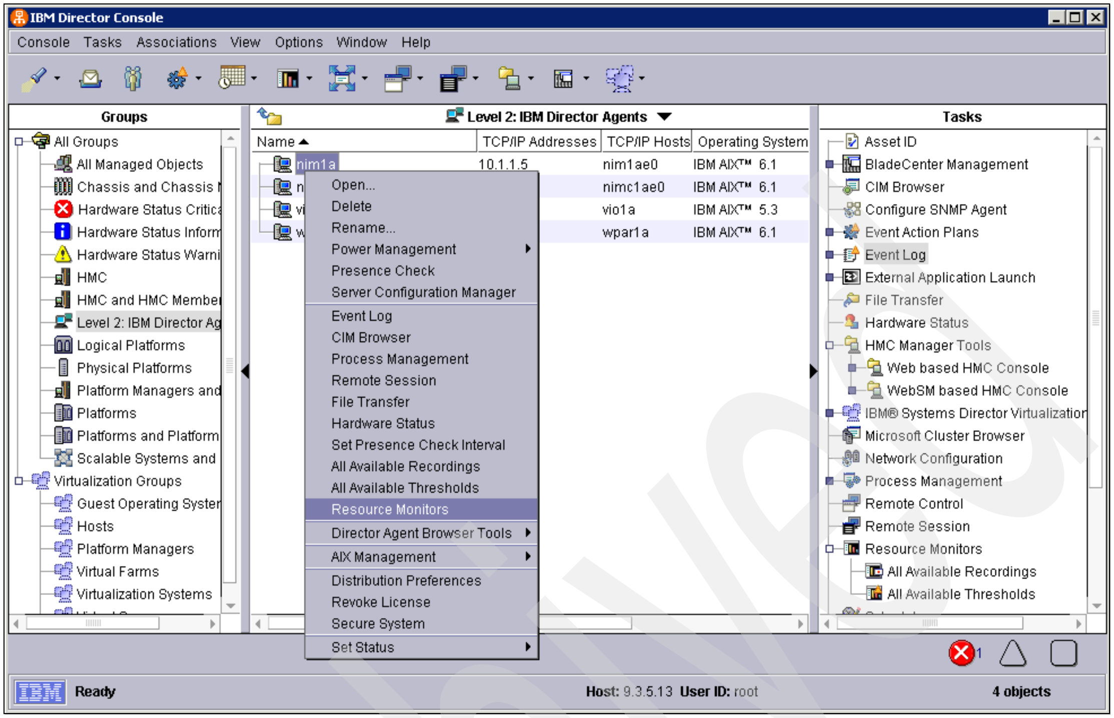

Name ▲	TCP/IP Addresses	TCP/IP Hosts	Operating System
nim1a	10.1.1.5	nim1ae0	IBM AIX™ 6.1
n		nimc1ae0	IBM AIX™ 6.1
vi		vio1a	IBM AIX™ 5.3
w		wpar1a	IBM AIX™ 6.1

Figure 4-40 How to open the Resource Monitors window

Figure 4-41 on page 222 shows an example of the available resources to be monitored for an AIX system with IBM Director Agent. On the right side of the window, several selected monitored metrics are shown. These numbers are periodically updated to reflect the real-time system workload.

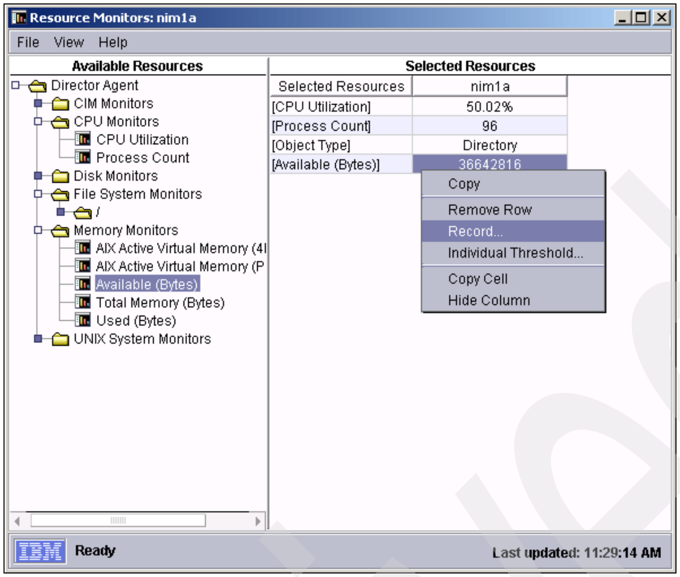

Selected Resources	nim1a
[CPU Utilization]	50.02%
[Process Count]	96
[Object Type]	Directory
[Available (Bytes)]	36642816

Last updated: 11:29:14 AM

Figure 4-41 Available monitored resources and the real -time resource metrics

IBM Director also offers a function to set a threshold for each performance metric. If the metric goes beyond the threshold, IBM Director triggers a predefined action. IBM Director also records the data. The recorded data can be shown in a table or graph. Figure 4-42 on page 223 shows an example of network utilization data. Unlike Tivoli Monitoring, IBM Director's grapher function is not customizable, and shows only one performance metric at a time.

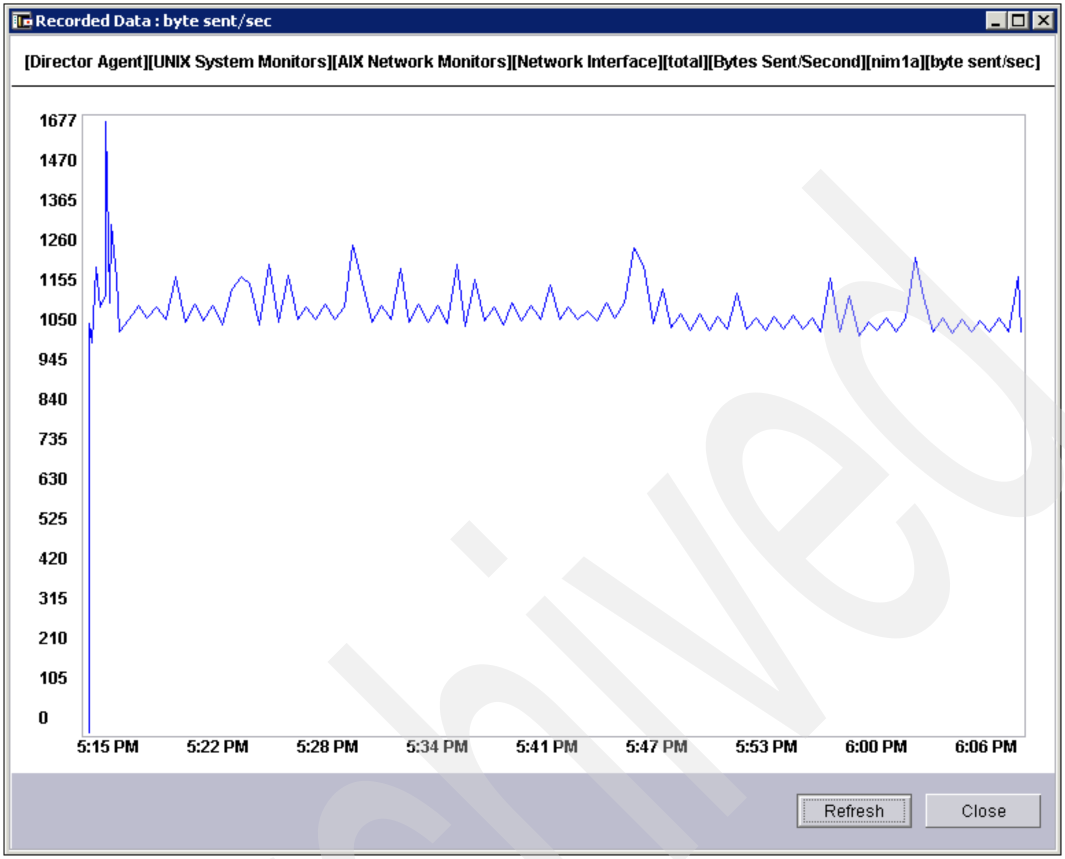
Recorded Data : byte sent/sec
[Director Agent][UNIX System Monitors][AIX Network Monitors][Network Interface][total][Bytes Sent/Second][nim1a][byte sent/sec]
1677
1470
1365
1260
1155
1050
945
840
735
630
525
420
315
210
105
0
5:15 PM
5:22 PM
5:28 PM
5:34 PM
5:41 PM
5:47 PM
5:53 PM
6:00 PM
6:06 PM
Refresh
Close

Figure 4-42 IBM Director - a grapher function

Figure 4-43 on page 224 shows an example of threshold configuration for a file system utilization metric. After the threshold is set, you can configure an automated action, which is activated when the resource metric goes beyond the threshold. With such an action plan, IBM Director can automatically execute any predefined operations, such as logging an event, sending an e-mail, or executing a program at the managed system. These operations can be used to notify a system administrator or recover from the undesirable situation.

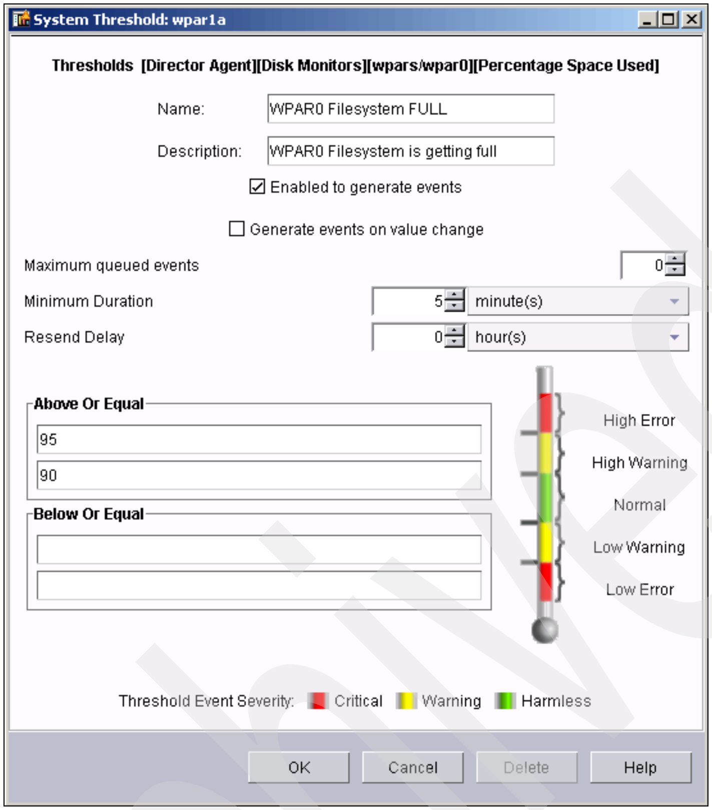
System Threshold: wpar1a

Thresholds [Director Agent][Disk Monitors][wpars/wpar0][Percentage Space Used]

Name: WPAR0 Filesystem FULL

Description: WPAR0 Filesystem is getting full

☑ Enabled to generate events

□ Generate events on value change

Maximum queued events: 0

Minimum Duration: 5 minute(s)

Resend Delay: 0 hour(s)

Above Or Equal: 95, 90

Below Or Equal:

Threshold Event Severity: 🟥 Critical 🟨 Warning 🟩 Harmless

Figure 4-43 Threshold configuration example

Figure 4-44 on page 225 shows an example of an event logged by IBM Director when a threshold is surpassed.

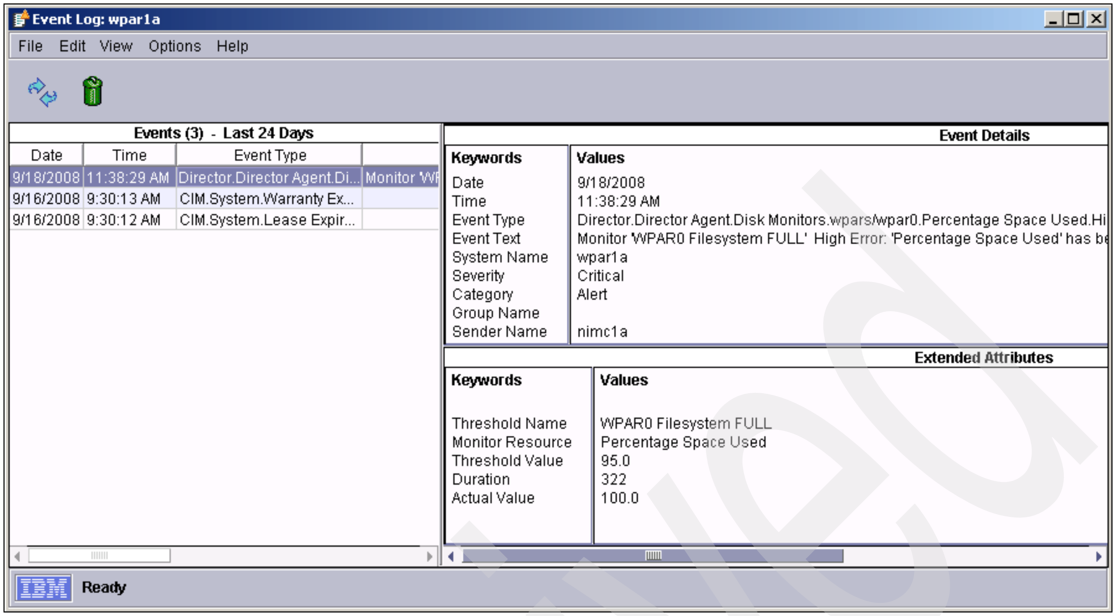

Events (3) - Last 24 Days

Date	Time	Event Type	
9/18/2008	11:38:29 AM	Director.Director Agent.Di...	Monitor 'WI
9/16/2008	9:30:13 AM	CIM.System.Warranty Ex...	
9/16/2008	9:30:12 AM	CIM.System.Lease Expir...	

Event Details

Keywords	Values
Date	9/18/2008
Time	11:38:29 AM
Event Type	Director.Director Agent.Disk Monitors.wpars/wpar0.Percentage Space Used.Hi
Event Text	Monitor 'WPAR0 Filesystem FULL' High Error: 'Percentage Space Used' has be
System Name	wpar1a
Severity	Critical
Category	Alert
Group Name	
Sender Name	nimc1a

Extended Attributes

Keywords	Values
Threshold Name	WPAR0 Filesystem FULL
Monitor Resource	Percentage Space Used
Threshold Value	95.0
Duration	322
Actual Value	100.0

Figure 4-44 Threshold event example

IBM Director firewall considerations

IBM Director processes require access to a number of ports in the network environment. See the following resources:

- Network protocols

  http://publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo_5.20/fqm0_r_network_protocols.html

- Ports used by IBM Director (Virtualization Engine port assignments)

  http://publib.boulder.ibm.com/infocenter/eserver/v1r2/index.jsp?topic=/veicinfo/eicarconcept_ports.htm

IBM Director monitored system requirements

IBM Director can manage Power Systems servers in two ways, with or without agent software installed. AIX servers without agent software are called Level-0 managed systems, and AIX servers with agent software are called Level-2 managed systems. Level-1 agent, available for some operating systems, is a subset of Level-2 agent. Level-1 agent is not available for AIX.

Level-2 managed systems provide the full IBM Director Agent functionality. The functionality of Level-2 IBM Director Agent on a managed system varies, depending on the operating system and platform. Level-0 managed systems provide a limited management functionality.

IBM Director publications

See the following resources for details about installing, configuring, and managing IBM Director:

- *IBM Director on System p5*, REDP-4219
- *Implementing IBM Director*, SG24-6188
- IBM Director V5.20 information center

  http://publib.boulder.ibm.com/infocenter/eserver/v1r2/index.jsp?topic=/diricinfo_5.20/fqm0_main.html
- Documentation and resources for IBM Director

  http://www.ibm.com/systems/management/director/resources/index.html

4.4.5 AIX built-in and third party commands

Several AIX built-in and third-party commands are useful to monitor both the overall system resource usage and the resource utilization of each tenant.

The topas command

The **topas** command gathers general information from different metrics on a partition. The major advantage of this tool is that you can see all of the important performance metrics of the partition at real-time in a single window.

Monitoring a single LPAR with topas

Example 4-1 shows an example of **topas** command real-time monitoring output. By default, this information is updated every two seconds.

Example 4-1 Real-time topas monitoring

```
Topas Monitor for host:    wpar1a               EVENTS/QUEUES    FILE/TTY
Mon Sep 15 17:21:46 2008   Interval:  2         Cswitch     176  Readch  1293.7K
                                                Syscall   82984  Writech     318
Kernel   14.7   |#####                       |  Reads     82799  Rawin         0
User      3.8   |##                          |  Writes        1  Ttyout      318
Wait      0.0   |                            |  Forks         0  Igets         0
Idle     81.5   |#######################     |  Execs         0  Namei         1
Physc =  0.04                     %Entc=  19.7  Runqueue    1.5  Dirblk        0
                                                Waitqueue   0.0
Network  KBPS   I-Pack  O-Pack   KB-In  KB-Out                   MEMORY
Total    35.3     49.0    33.0    29.1     6.2  PAGING           Real,MB    2048
                                                Faults        1  % Comp     37.7
```

```
Disk    Busy%     KBPS     TPS KB-Read KB-Writ  Steals        0  % Noncomp   7.7
Total     7.5    129.7    30.4     0.0   129.7  PgspIn        0  % Client    7.7
                                                PgspOut       0
FileSystem        KBPS     TPS KB-Read KB-Writ  PageIn        0  PAGING SPACE
Total              0.0     0.0     0.0     0.0  PageOut       4  Size,MB     512
                                                Sios          4  % Used      2.2
WLM-Class (Active)     CPU%    Mem%  Blk-I/O%                    % Free     98.8
System                    2      16         2   NFS (calls/sec)
Unmanaged                 0      29         0   SerV2         0  WPAR Activ    5
                                                CliV2         0  WPAR Total    5
Name            PID CPU% PgSp Class             SerV3         0  Press: "h"-help
od           544798 29.2  0.2 wpar0             CliV3         0         "q"-quit
topas        655526  0.6  2.5 System
topas        798948  0.6  2.4 System
xmtopas      815282  0.4  0.6 System
wlmsched      65568  0.3  0.5 System
```

Monitoring multiple LPARs with topas

You can obtain inter-partition reports with the **topas -C** command (called the Cross-Partition view). Example 4-2 shows an example of cross-LPAR real-time monitoring by topas.

Example 4-2 Cross-Partition view

```
Topas CEC Monitor             Interval:  10            Mon Sep 15 17:28:00 2008
Partitions Memory (GB)           Processors
Shr:  7    Mon:18.9  InUse:16.4  Shr:2.2  PSz:  3   Don: 0.0 Shr_PhysB  0.35
Ded:  1    Avl:   -              Ded:  1  APP:  2.6 Stl: 0.0 Ded_PhysB  0.08

Host         OS  M Mem InU Lp  Us Sy Wa Id  PhysB  Vcsw Ent  %EntC PhI
-------------------------------------shared-------------------------------------
nimc1a       A61 U 2.0 2.0  2  77 20  0  1   0.25  621  0.20 123.8   6
uam1a        A61 U 2.0 1.3  8   1  1  0 96   0.04  919  1.00   3.8   1
wpar1a       A61 U 2.0 0.9  2   2 11  0 86   0.03  314  0.20  15.0   1
vio1a        A53 U 2.0 1.5  2   0  4  0 95   0.01  669  0.20   7.3   0
tpm1a        A53 U 2.9 2.8  2   1  3  0 94   0.01  157  0.20   6.0   1
nim1a        A61 U 2.0 2.0  2   0  2  0 97   0.01  236  0.20   3.1   0
mozuku0      A53 U 2.0 2.0  2   0  1  0 98   0.00  135  0.20   1.7   0

Host         OS  M Mem InU Lp  Us Sy Wa Id  PhysB  Vcsw  %istl %bstl
------------------------------------dedicated-----------------------------------
tsm1a        A61 S 4.0 3.9  2   6  2  0 90   0.08  534    0.00  0.00
```

For details about the **topas** command, see “13.2 Monitoring real time consumption” and “13.3 Monitoring cross-partition real time consumption” in *PowerVM Virtualization on IBM System p: Managing and Monitoring,* SG24-7590

Considerations for using topas command

For cross-LPAR monitoring feature to work, the **`topas -C`** command requires that the xmtopas daemon run on the monitored systems. The **`topas`** command also is able to monitor the Virtual I/O Server. In that case, you have to edit the `/etc/inetd.conf` file to make sure the **`xmtopas`** command initiates at boot time. The daemon is disabled by default. In addition, the network connection between the monitoring system (running **`topas -C`**) and the monitored system (running **`xmtopas`**) is required, as shown in Table 4-6.

Table 4-6 Required network connection for topas Cross-Partition feature

Service	Source	Destination	SRC port	DST port
xmquery	Monitoring system running **`topas -C`** command	Monitored system running xmtopas daemon	Any	2279:tcp, 2279:udp

Note that the **`topas`** command uses broadcast packets to find other LPARs. Further information is available in the IBM Systems Information Center. See how to add a host to the topas external subnet search file (Rsi.hosts):

http://publib.boulder.ibm.com/infocenter/systems/index.jsp?topic=/com.ibm.aix.prftungd/doc/prftungd/adding_rsihosts.htm

The nmon tool

Although nmon tool is not an officially supported tool, it is widely used and very helpful to display, collect, and analyze performance data of a Power Systems partition which is running AIX or Linux. The nmon tool can display the real-time performance data such as topas, or it can the collect performance data into a file with a specified interval, for a certain period of time. A reporting tool called nmon2rrd can be very useful, it converts the nmon output file into Web pages automatically. Another reporting tool called nmonanalyser is also very helpful, it can convert an nmon output file to graphical charts by using a Microsoft Excel® macro. A reporting tool for topas, the **`topasout`** command, supports the nmonanalyser tool, it can generate the input file for the nmonanalyser.

For more information, see the following resources:

- The nmon tool for AIX & Linux Performance Monitoring

  http://www.ibm.com/developerworks/wikis/display/WikiPtype/nmon

- AIX tips for the nmon2rrd tool

  http://www.aixtips.com/AIXtip/nmon2rrdv1.htm

- The nmonanalyser tool

  http://www.ibm.com/developerworks/wikis/display/WikiPtype/nmonanalyser

- The **topasout** command

  http://publib.boulder.ibm.com/infocenter/systems/index.jsp?topic=/com.ibm.aix.cmds/doc/aixcmds5/topasout.htm

5

Accounting

In this chapter we discuss accounting and chargeback in a multitenant AIX environment. We start with an overview of accounting and chargeback, describe the tools available from IBM, and describe an actual implementation by using these tools.

This chapter contains the following topics:

5.1 Introduction to financial management

Financial management is a key concept of any IT operation. The three primary activities for financial management are:

- Budgeting
- Accounting
- Chargeback

If we consider the internal IT department of a company as a service provider, the department often operates with a cost-center model, where the IT department charges the other departments of the company based on resource usage. In this case, the total chargeback is close to the total operational costs. A service provider that sells services to other companies obviously has to add its gross profit margin on top of the operational costs.

In this chapter, we concentrate on the accounting and chargeback activities, how they relate to multitenant AIX environments, and the IBM tools that are available to help implement them.

5.2 Accounting and chargeback considerations

In this section, we discuss the following two topics:

- Accounting, which is used for monitoring and controlling operational costs, and can be used for capacity management and budgeting expenses
- Chargeback, which is a general term that can be used to represent either charging of resource usage to different business units of the same company, or billing other companies based on their resource usage

Accounting is usually based on the actual resource usage measurements. Chargeback, however, can be based on the actual resource usage, several completely different criteria, or both, depending on your business model. As an example, the chargeback for a system running SAP R/3® can be based on the number of named users within the SAP R/3 application, instead of actual resource consumption measurements.

5.2.1 Accounting policy

Chargeback is usually based on accounting information. Accounting reflects resource usage measurements, but can also include other information, such as number of user accounts, power consumption, floor space, or fixed amount of money.

Before we start collecting accounting data, we consider what is format of data for accounting. For example, if we want resource usage as the accounting data, we consider the frequency and the overhead of collecting data.

In a multitenant environment, the accounting data is located in a tenant's environment, therefore, the risk of tampering of the accounting data exists. The service provider has to develop safeguards against tampering or corruption of the accounting data.

Generally, consider the following resource usage measurements for accounting:

- CPU usage
- Memory usage
- Disk, I/O usage
- File system consumption

For a more sophisticated accounting, we must compromise between data granularity and accurate data.

5.2.2 Chargeback policy

The chargeback policy is a major part of the business model for any service provider, and it can be very different for each provider. Therefore we discuss only considerations that must addressed.

When creating the chargeback policy, a service provider has to consider how much detail about their operational costs they are willing to reveal to their tenants. Providing highly granular accounting information tends to lead to commoditization of the service and gives a competitive advantage to other service providers.

In the order of significance, the chargeback policy generally has to be:

1. Fair

   The policy should be based on criteria that makes sense to the tenant, such as resources allocated to the tenant or resources actually used by the tenant.

2. Simple

   The number of different resources used for chargeback purposes should be kept to a minimum.

3. Accurate

   The data used for chargeback needs to be as accurate as possible. It also needs to be archived so that it can be reviewed later, in case of disputes over charges.

5.3 AIX Advanced Accounting overview

Advanced Accounting is a feature that was introduced in AIX5L V5.3, in addition to the traditional UNIX accounting. It is based on mainframe technology, such as interval accounting and transaction accounting, and provides much more detailed accounting data than is traditionally available in UNIX.

Advanced Accounting supports the new virtualization technologies available with the Power Systems and AIX, such as logical partitions (LPARs), micro-partitioning, and workload partitions (WPARs). Advanced accounting can also be configured and used in the Virtual I/O Server.

Advanced Accounting provides usage information for a wide variety of system resources, and enables you to develop comprehensive chargeback strategies. The data sources include resources such as disks, network interfaces, virtual devices, file systems, processors, and memory.

The details of configuring AIX Advanced Accounting are discussed in:

- Chapter 5 “Advanced Accounting,” in *Accounting and Auditing on AIX 5L,* SG24-6396.
- 5.6, “Configuring data collection in AIX” on page 241.

5.4 Tivoli Usage and Accounting Manager overview

Tivoli Usage and Accounting Manager is a general purpose tool that performs the following tasks:

- Collects resource usage data
- Assigns account codes to each resource
- Provides chargeback rates for each resource

In addition, the Tivoli Usage and Accounting Manager provides reports, such as the Usage Trend reports, which can be used as input for capacity management.

The Tivoli Usage and Accounting Manager supports data collectors for several operating systems, hardware platforms, and applications. It also provides a universal collector for including resource utilization data contained in any usage log, spreadsheet, or database table. To minimize the resources required for data collection, Tivoli Usage and Accounting Manager collects data by using the native accounting tools of each environment. In this chapter, we focus on capturing the resource utilization data from the AIX Advanced Accounting feature, and using that data for reporting purposes.

For more information about installing, configuring, and using Tivoli Usage and Accounting Manager, see *IBM Tivoli Usage and Accounting Manager V7.1 Handbook,* SG24-7404. Also see the product documentation for Tivoli Usage and Accounting Manager:

```
http://publib.boulder.ibm.com/infocenter/tivihelp/v3r1/index.jsp?topic=
/com.ibm.ituam.doc_7.1/welcome.htm
```

5.4.1 Tivoli Usage and Accounting Manager example environment

Tivoli Usage and Accounting Manager consists of four components:

- Database server
- Administration server
- Processing server
- Reporting server

These components can all run on a single server, or they can be divided into different servers.

The first three components can be run on either UNIX or Windows platforms. However, because of the Microsoft software products used, the reporting server component can only be run on Windows 2003.

Our Tivoli Usage and Accounting Manager test environment has two servers:

- AIX 6.1 server uam1a
  - DB2 Universal Database™ for AIX Version 9.5
  - Embedded WebSphere Application Server 6.1
  - Integrated Solutions Console
  - Tivoli Usage and Accounting Manager processing engine
  - Tivoli Usage and Accounting Manager data collectors
- Windows 2003 server uamwin
  - Microsoft Report Viewer
  - Microsoft Internet Information Services
  - Tivoli Usage and Accounting Manager reporting application

In addition, every AIX partition and every Virtual I/O Server has the TUAM AIX Advanced Accounting data collector installed. The environment is shown in Figure 5-1.

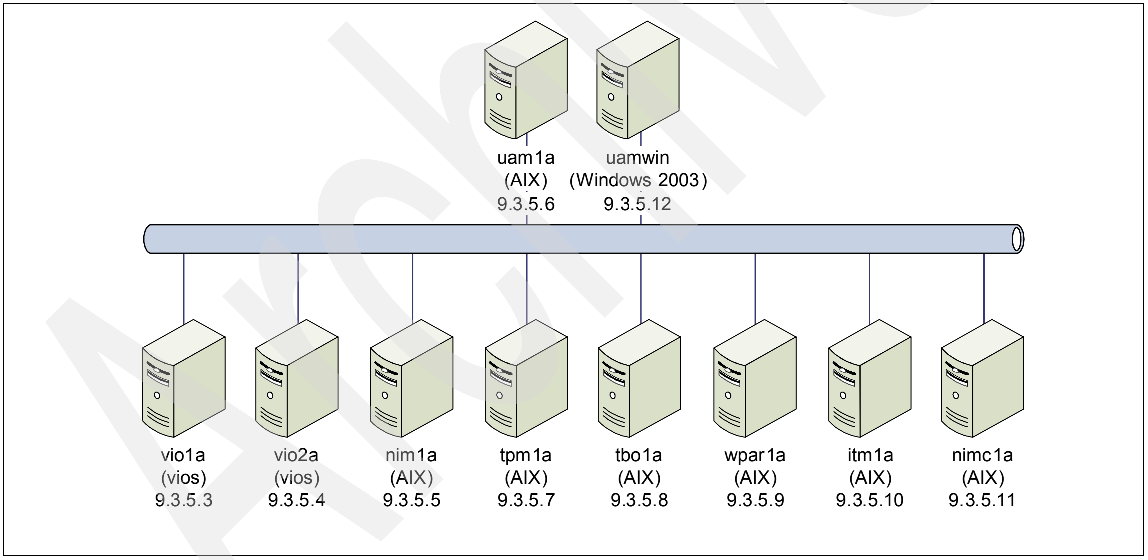

Figure 5-1 Tivoli Usage and Accounting Manager example environment

Note: We do not have firewalls in our test environment. In any real production environment, firewalls are usually required between the tenant networks and the management network.

5.5 Account codes and rate tables

Several basic configuration steps have to be completed in the Tivoli Usage and Accounting Manager server before any data collection can be done:

- Defining the account code structure
- Setting up the CPU normalization
- Mapping host names to account codes
- Setting up the rate tables
- Defining clients

5.5.1 Defining the account code structure

The account code is the primary identifier that signifies who should be billed for specific system usage. The account code consists of several fixed-length fields.

The account code structure has to be defined before any data collection or processing. Because the account code is used as the primary key for all the collected accounting data, changing the account code structure can make all the collected data invalid.

In our example, we define the account code structure as described in Table 5-1. Each customer can have multiple contract numbers, which can each contain multiple applications, and each application can contain multiple hosts.

Table 5-1 Sample account code structure

Description	Length
Customer	8
Contract number	12
Application	8
Host	32

The account code structure can be changed in the Integrated Solutions Console, by selecting **Usage and Accounting Manager** → **System Maintenance** → **Account Code Structure**, as shown in Figure 5-2 on page 238.

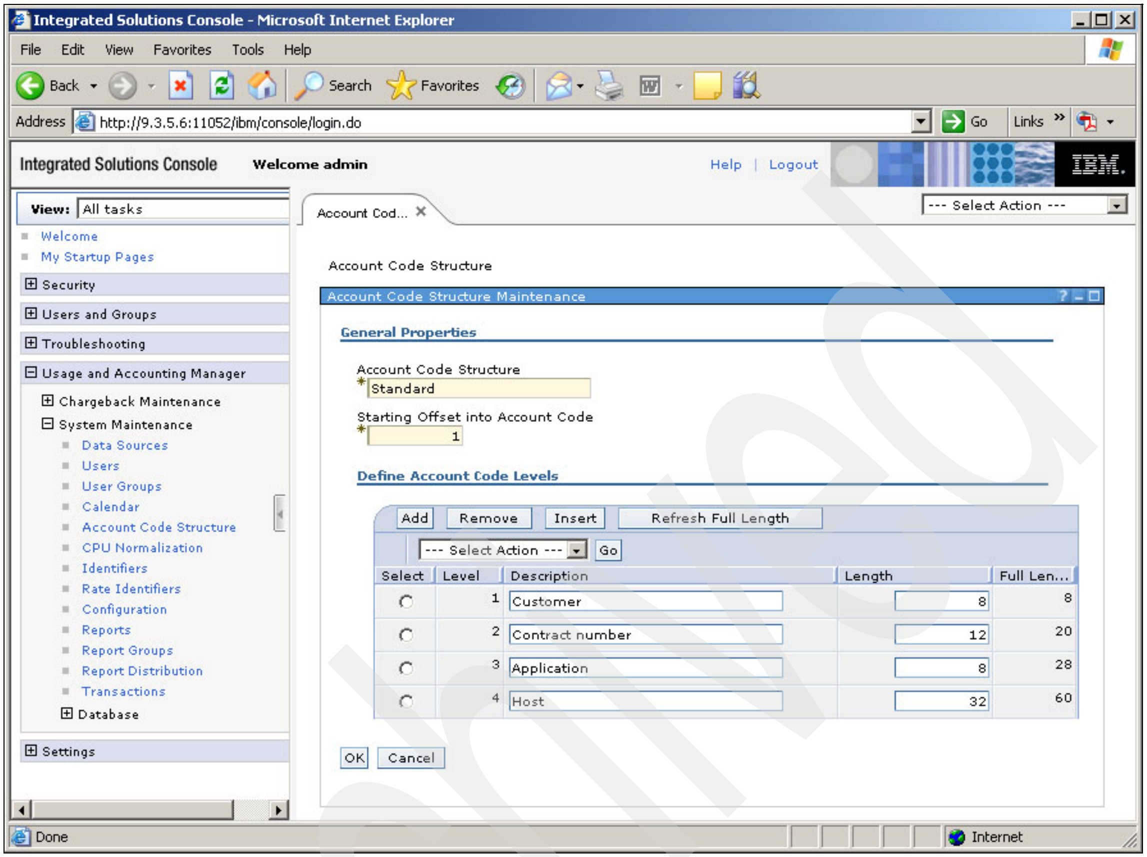

Account Code Structure

Account Code Structure Maintenance

General Properties

Account Code Structure
* Standard

Starting Offset into Account Code
* 1

Define Account Code Levels

Select	Level	Description	Length	Full Len...
○	1	Customer	8	8
○	2	Contract number	12	20
○	3	Application	8	28
○	4	Host	32	60

Figure 5-2 Tivoli Usage and Accounting Manager account code structure

5.5.2 Setting up the CPU normalization

The performance provided by a single processor can vary significantly based on several variables, such as:

- Processor generation
- Clock speed
- Memory architecture

TUAM support the changing of CPU normalization settings. Doing so is not required, and generally not recommended. If you have to account for different types of processors, and have comparable values for billing purposes, the metering values for different hardware can be normalized. Those values can then be set in the Integrated Solutions Console by selecting **Usage and Accounting Manager** → **System Maintenance** → **CPU Normalization**.

In our example, we chose not to implement CPU normalization because we are using only one CPU model.

5.5.3 Mapping host names to account codes

Several ways are available to create the account code information from the host names. We chose to implement a simple comma-separated lookup table that defines the account code structure. The table is named `accttabl.txt`, located in `/opt/ibm/tuam/processes`, and shown in Example 5-1.

Example 5-1 accttabl.txt

```
vio1a,,Cust1   000000000001VIOS
vio2a,,Cust1   000000000001VIOS
nim1a,,Cust1   000000000001Appl1
uam1a,,Cust1   000000000001Appl2
tpm1a,,Sarah   000000000002TPM
tbo1a,,Sarah   000000000002TPM
itm1a,,Carlos  000000000003ITM
wpar1a,,Masahiko000000000004WPAR
nimc1a,,Carlos  000000000003ITM
```

The first and second fields define the range of host names. A single line in this file can represent multiple hosts, as long as the start of the host name matches the first and second fields. This is the reason why we chose to have a table separate from the Tivoli Usage and Accounting Manager hosts file used in data collection in 5.7, “Collecting and loading AIX accounting data” on page 247. The third field represents the first three fields (or 28 characters) of the account code, and the host name is automatically appended to it.

5.5.4 Setting up the rate tables

The rate tables define the charge rates for all input metrics. The basic installation of Tivoli Usage and Accounting Manager includes a rate table named STANDARD, but you can define multiple rate tables, if required. The rate table can be chosen separately for each tenant, based on the account code.

The rate tables can be edited in the Integrated Solutions Console by selecting **Usage and Accounting Manager** → **Chargeback Maintenance** → **Rates**. All AIX Advanced Accounting rate names begin with AAID so that the names are easier to find in the Integrated Solutions Console. Several entries of the rate table are shown in Figure 5-3 on page 240.

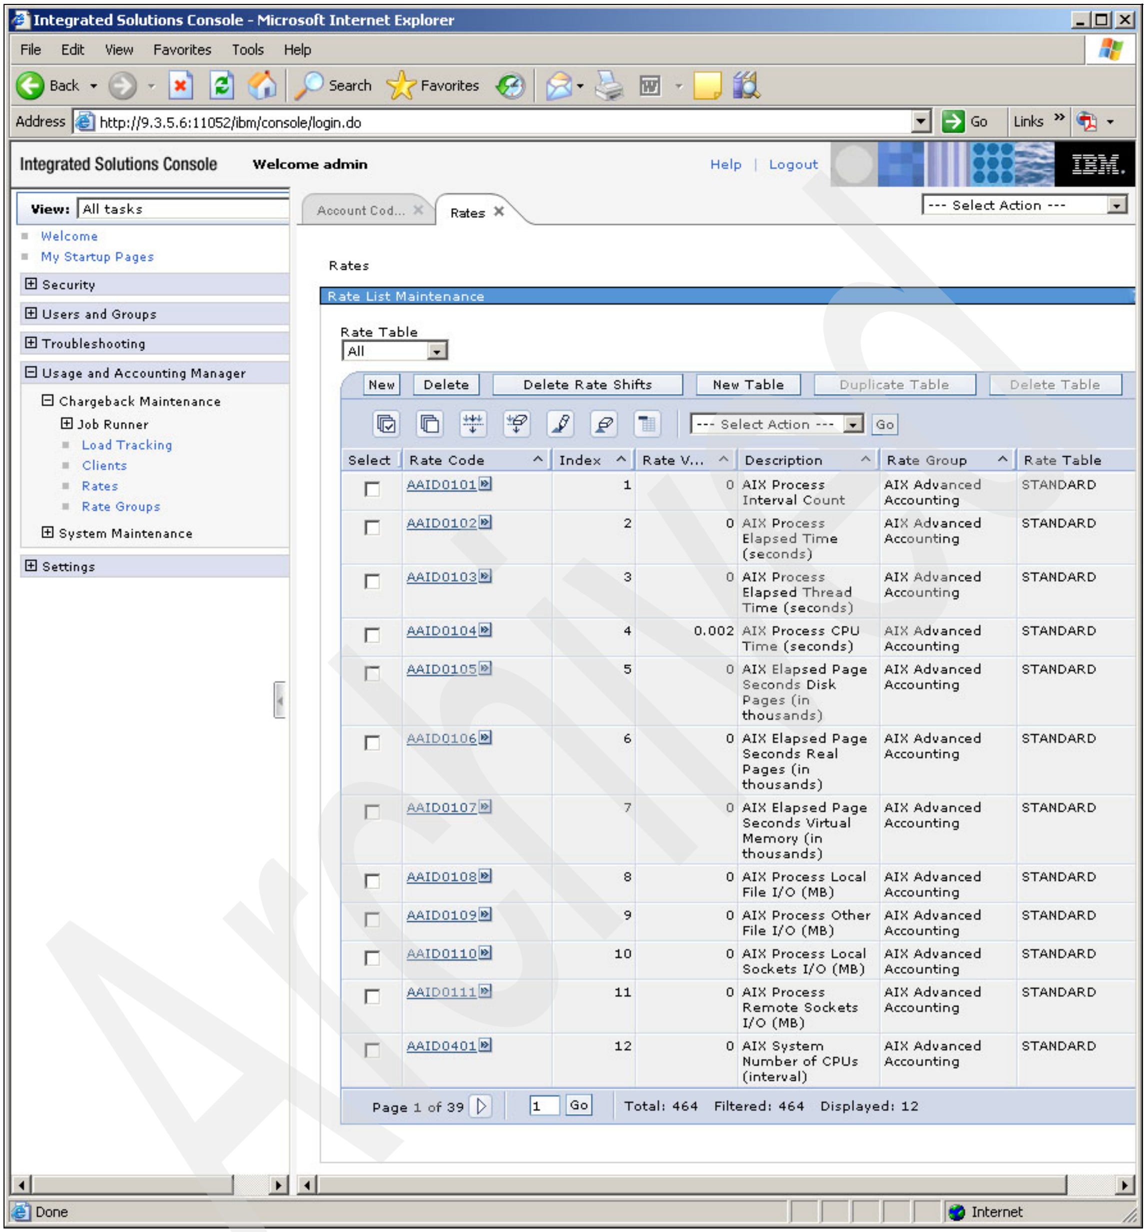

Rate Table: All

Select	Rate Code	Index	Rate V...	Description	Rate Group	Rate Table
□	AAID0101	1	0	AIX Process Interval Count	AIX Advanced Accounting	STANDARD
□	AAID0102	2	0	AIX Process Elapsed Time (seconds)	AIX Advanced Accounting	STANDARD
□	AAID0103	3	0	AIX Process Elapsed Thread Time (seconds)	AIX Advanced Accounting	STANDARD
□	AAID0104	4	0.002	AIX Process CPU Time (seconds)	AIX Advanced Accounting	STANDARD
□	AAID0105	5	0	AIX Elapsed Page Seconds Disk Pages (in thousands)	AIX Advanced Accounting	STANDARD
□	AAID0106	6	0	AIX Elapsed Page Seconds Real Pages (in thousands)	AIX Advanced Accounting	STANDARD
□	AAID0107	7	0	AIX Elapsed Page Seconds Virtual Memory (in thousands)	AIX Advanced Accounting	STANDARD
□	AAID0108	8	0	AIX Process Local File I/O (MB)	AIX Advanced Accounting	STANDARD
□	AAID0109	9	0	AIX Process Other File I/O (MB)	AIX Advanced Accounting	STANDARD
□	AAID0110	10	0	AIX Process Local Sockets I/O (MB)	AIX Advanced Accounting	STANDARD
□	AAID0111	11	0	AIX Process Remote Sockets I/O (MB)	AIX Advanced Accounting	STANDARD
□	AAID0401	12	0	AIX System Number of CPUs (interval)	AIX Advanced Accounting	STANDARD

Page 1 of 39 ▷ 1 Go Total: 464 Filtered: 464 Displayed: 12

Figure 5-3 Tivoli Usage and Accounting Manager rate tables

5.5.5 Defining clients

In Tivoli Usage and Accounting Manager, a client is defined by the start of the account code described previously. In most cases, a recommendation is define a separate client for each tenant.

Although defining clients is optional, it is good practice because it enables you to:

- Specify the contact information of the tenants, such as mailing address and telephone numbers.
- Use different rate tables for different servers.
- Give tenants limited access to reports for their servers, as described in 5.9, “Providing resource utilization view to tenants” on page 266.

The clients can be defined in the Integrated Solutions Console by selecting **Usage and Accounting Manager** → **Chargeback Maintenance** → **Clients**.

The definition of new clients as well as their usage in a multitenant environment are discussed in more detail in 5.9.1, “Defining clients” on page 266.

5.6 Configuring data collection in AIX

Configuration of data collection has to be automated as part of the provisioning process. However, the actual required steps are described in this section.

5.6.1 Overview of the data collection on tenant side

In our implementation, data collection on the tenant side consists of three parts:

- AIX Advanced Accounting writes the accounting information into predefined files in the directory `/var/aacct`.
- Every hour, any new primary accounting files are copied from the directory `/var/aacct` to the directory `$TUAM/history`, and released to be reused.
- Every night, the files in the directory `$TUAM/history` are processed to create separate accounting files for each Advanced Accounting record type in `$TUAM/CS_input_source`.

The Tivoli Usage and Accounting Manager server collects the separate accounting files daily by using the **`scp`** command.

The data flow is described in Figure 5-4 on page 242.

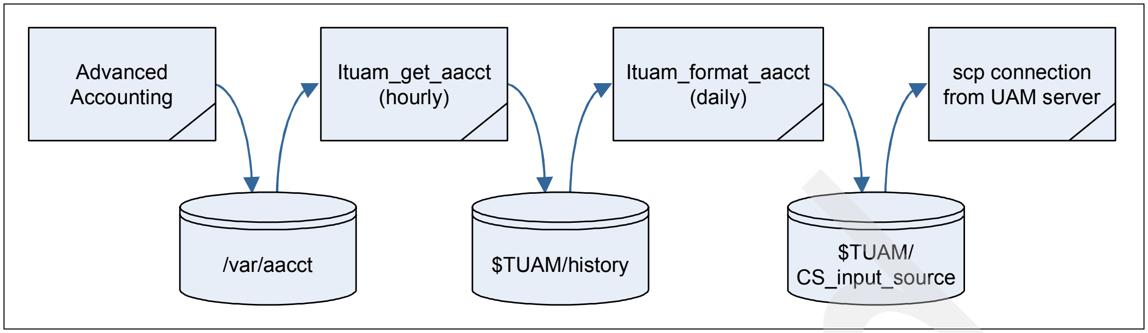

Figure 5-4 Tivoli Usage and Accounting Manager data flow on the tenant side

In the figure, $TUAM represents our Tivoli Usage and Accounting Manager data collector installation directory, `/opt/ibm/tuam/collectors/unix`.

5.6.2 Installation prerequisites

For installing the AIX Advanced Accounting and Tivoli Usage and Accounting Manager data collector, we created a shared installation directory, which contains the files described in Table 5-2.

Table 5-2 Files in installation directory

Filename	Description
tuam_unpack_uc_collector	Installation script for the Tivoli Usage and Accounting Manager package
ituam_uc_aix5.tar	Tivoli Usage and Accounting Manager data collector package for AIX
ituam_schedule.sh	Our custom script to manage hourly and daily processing of AIX Advanced Accounting data
install-tuam.sh	Our custom installation script for setting up AIX Advanced Accounting and Tivoli Usage and Accounting Manager
id_dsa.pub	SSH public key for the root user of our Tivoli Usage and Accounting Manager server, required for setting up data collection, generated with the **`ssh-keygen`** command

In addition, the openssh packages have to be installed on the tenant system. The installation script is shown in Example 5-2 on page 243, and the hourly data processing script is shown in Example 5-3 on page 244.

> **Note:** Using the root user ID for non-administrative tasks, such as data transfer is not good practice. Therefore we create a user ID with no root privileges for the Tivoli Usage and Accounting Manager data transfer.

Example 5-2 Installation script: install-tuam.sh

```
#!/usr/bin/ksh

INSTPATH=/opt/ibm/tuam/collectors/unix
TUAMUSER=tuam

# First set up AIX Advanced Accounting

acctctl on
acctctl fadd /var/aacct/aacct0.dat 1
acctctl fadd /var/aacct/aacct1.dat 1
acctctl fadd /var/aacct/aacct2.dat 1
acctctl fadd /var/aacct/aacct3.dat 1
acctctl fadd /var/aacct/aacct4.dat 1
acctctl isystem 5
acctctl iprocess 5
acctctl agproc on
acctctl agke on
acctctl agarm on
mkitab 'aacct:2:once:/usr/bin/acctctl on >/dev/console 2>&1'

# Next set up userid for TUAM collector and enable ssh from the TUAM server
mkgroup ${TUAMUSER}
mkuser pgrp=${TUAMUSER} ${TUAMUSER}
mkdir ~${TUAMUSER}/.ssh
cp ./id_dsa.pub ~${TUAMUSER}/.ssh/authorized_keys2
chown -R ${TUAMUSER}.${TUAMUSER} ~${TUAMUSER}/.ssh
chmod 755 ${TUAMUSER} ${TUAMUSER}/.ssh

# Set up the script for data collection
mkdir -p /usr/local/bin
cp ./ituam_schedule.sh /usr/local/bin

# Install the TUAM collector
mkdir -p ${INSTPATH}
cp ./tuam_unpack_uc_collector ./ituam_uc_aix5ar ${INSTPATH}
cd ${INSTPATH}
# use aacct_config=true when calling tuam_unpack_uc_collector to prevent
# traditional UNIX process accounting from being started
./tuam_unpack_uc_collector path=${INSTPATH} user=${TUAMUSER} aacct_config=true
rm ${INSTPATH}/tuam_unpack_uc_collector ${INSTPATH}/ituam_uc_aix5.tar
```

```
# Adjust TUAM AAA collector to include client vscsi
cat ${INSTPATH}/data/A_config.par|sed
'/^AACCT_TRANS_IDS/s/^.*$/AACCT_TRANS_IDS="1,4,6,7,8,11"/' >/tmp/A_config.tmp
cp /tmp/A_config.tmp ${INSTPATH}/data/A_config.par
rm /tmp/A_config.tmp

# Adjust crontab for the data collection
crontab -l|sed '/.*aacct/s/^/#/' >/tmp/tempcrontab
echo '59 * * * * /usr/local/bin/ituam_schedule.sh
>/opt/ibm/tuam/collectors/unix/log/ituam_schedule.log 2>&1' >>/tmp/tempcrontab
crontab /tmp/tempcrontab
rm /tmp/tempcrontab
```

Example 5-3 Hourly data processing script: ituam-schedule.sh

```
#!/usr/bin/ksh

ITUAMHOME=/opt/ibm/tuam/collectors/unix/scripts/aacct
ITUAMGET=${ITUAMHOME}/ituam_get_aacct
ITUAMFORMAT=${ITUAMHOME}/ituam_format_aacct

# get current date
START=$(date +%Y%m%d)

# call ITUAMGET
${ITUAMGET}

# wait for 2 minutes
sleep 120

#check if date has changed
if [ ${START} -eq $(date +%Y%m%d) ]; then
        echo same date - no preprocessing needed
else
        # date changed, run format script for previous day
        ${ITUAMFORMAT} ${START}
fi
```

5.6.3 Installation process

The installation process for the partition tpm1a, using the scripts in the previous examples, is shown in Example 5-4 on page 245.

Example 5-4   Installation process of Tivoli Usage and Accounting Manager data collectors

```
root@tpm1a:/mnt/tuam-install $ ls -l
total 42248
-rw-r--r--    1 root     system          600 Aug 28 11:53 id_dsa.pub
-rwxr-xr-x    1 root     system         1599 Sep  2 09:02 install-tuam.sh
-rwxr-xr-x    1 root     system          456 Aug 28 13:31 ituam_schedule.sh
-rw-r--r--    1 root     system     21575680 Aug 28 11:43 ituam_uc_aix5.tar
-rwxr-xr-x    1 root     system        36719 Aug 28 11:43 tuam_unpack_uc_collector
root@tpm1a:/mnt/tuam-install $ ./install-tuam.sh
./tuam_unpack_uc_collector: Begin ITUAM UNIX/Linux Data Collector Installation

./tuam_unpack_uc_collector: ITUAM UNIX/Linux Data Collector will be installed in
/opt/ibm/tuam/collectors/unix
    Nodename                            : tpm1a
    Platform Type                       : AIX
    Distribution tar file               : ituam_uc_aix5.tar
    ITUAM UNIX/Linux Collector Home     : /opt/ibm/tuam/collectors/unix
    ITUAM UNIX/Linux Collector User     : tuam
    ITUAM UNIX/Linux cs_method          : HOLD
    ITUAM UNIX/Linux server             :
    ITUAM UNIX/Linux cs_user            :
    Config AACCT                        : FALSE

./tuam_unpack_uc_collector: Unpacking ituam_uc_aix5.tar in /opt/ibm/tuam/collectors/unix
x accounting
x accounting/README.txt, 0 bytes, 0 tape blocks
...
x scripts/oracle/ituam_view.sql, 3001 bytes, 6 tape blocks

./tuam_unpack_uc_collector: Initialize /opt/ibm/tuam/collectors/unix/accounting directory
./tuam_unpack_uc_collector: Initialize /opt/ibm/tuam/collectors/unix/bin directory
./tuam_unpack_uc_collector: Initialize /opt/ibm/tuam/collectors/unix/etc directory
./tuam_unpack_uc_collector: Initialize /opt/ibm/tuam/collectors/unix/scripts directory
./tuam_unpack_uc_collector: Initialize /opt/ibm/tuam/collectors/unix/data directory
./tuam_unpack_uc_collector: Initializing ITUAM Configuration File.
    Setting - ITUAM_ACCDAT
    Setting - ITUAM_BIN
    Setting - ITUAM_DATA
    Setting - ITUAM_DESCR
    Setting - ITUAM_ETC
    Setting - ITUAM_EXAMPLES
    Setting - ITUAM_HELP
    Setting - ITUAM_HISTORY
    Setting - ITUAM_HOME
    Setting - ITUAM_LOG
    Setting - ITUAM_SCRIPTS
    Setting - ITUAM_SUPER
    Setting - ITUAM_USER
    Setting - ITUAM_SERVER
    Setting - ITUAM_DEST
    Setting - ITUAM_ACCOUNT
    Setting - ITUAM_GROUP
    Setting - CS_PLATFORM
    Setting - CS_USER
    Setting - CS_KEY
    Setting - CS_UPATH
    Setting - CS_METHOD
    Setting - CS_PROC_PATH

********************************************************************************
Starting ITUAM/UNIX create_A_storage.par Script at Tue Sep  2 10:00:01 CDT 2008
********************************************************************************

ITUAM/UNIX create_A_storage.par: Moving /opt/ibm/tuam/collectors/unix/data/A_storage.par to
/opt/ibm/tuam/collectors/unix/data/A_storage.par.bck
ITUAM/UNIX create_A_storage.par: Creating new /opt/ibm/tuam/collectors/unix/data/A_storage.par

********************************************************************************
Ending ITUAM/UNIX create_A_storage.par Script at Tue Sep  2 10:00:01 CDT 2008
********************************************************************************

./tuam_unpack_uc_collector: Initialize /opt/ibm/tuam/collectors/unix/description directory
./tuam_unpack_uc_collector: Initialize /opt/ibm/tuam/collectors/unix/examples directory
./tuam_unpack_uc_collector: Initialize /opt/ibm/tuam/collectors/unix/help directory
./tuam_unpack_uc_collector: Updating root crontab
./tuam_unpack_uc_collector: Start UNIX/Linux Process Accounting

********************************************************************************
Starting ITUAM/UNIX turnacct Script at Tue Sep  2 10:00:03 CDT 2008
********************************************************************************

ITUAM/UNIX turnacct: Turning accounting on ...

********************************************************************************
Ending ITUAM/UNIX turnacct Script at Tue Sep  2 10:00:03 CDT 2008
********************************************************************************

./tuam_unpack_uc_collector: ITUAM UNIX/Linux Data Collector installation complete
root@tpm1a:/mnt/tuam-install $
```

5.7 Collecting and loading AIX accounting data

In the previous section, we installed the necessary tools to create the input data for the Tivoli Usage and Accounting Manager into the tenant systems. The next step is to collect the data from all tenant systems into our Tivoli Usage and Accounting Manager server and load the data into the Tivoli Usage and Accounting Manager database.

5.7.1 Collecting AIX accounting data from the tenant systems

The two ways to collect the data are:

- Push-mode, where the tenant systems are responsible for copying the data to the Tivoli Usage and Accounting Manager server
- Pull-mode, where the Tivoli Usage and Accounting Manager server is responsible for copying the data from all the tenant systems

For security reasons, we do not want to allow any network connections from the tenant systems to our Tivoli Usage and Accounting Manager server. Therefore, we choose to implement the pull-mode method. We also choose to implement the actual transfer with scp.

5.7.2 Loading AIX accounting data into the database

The preprocessing of the AIX accounting data and loading it to the Tivoli Usage and Accounting Manager database is performed by Tivoli Usage and Accounting Manager jobs.

We copied the sample job description XML file, `SampleAIXAA.xml`, from the Tivoli Usage and Accounting Manager sample job files directory to the job files directory to work as a template. In this template file, we commented out the lines that are specifying the input files that we do not collect from the AIX servers. We also added the necessary steps for account code lookup to implement the account code structure we defined in 5.5.1, “Defining the account code structure” on page 237. The modified template job file, with the account code lookup steps highlighted, is shown in Example 5-5 on page 248.

Example 5-5 SampleAIXAA.xml

```
<?xml version="1.0" encoding="utf-8"?>
<!--
 *************************************************************** {COPYRIGHT-TOP}
 * Licensed Materials - Property of IBM
 * IBM Tivoli Usage and Accounting Manager
 * 5724-O33, 5765-UAV, 5765-UA7, 44E7863
 * (c) Copyright IBM Corp. 2004, 2007
 *
 * The source code for this program is not published or otherwise
 * divested of its trade secrets, irrespective of what has been
 * deposited with the U.S. Copyright Office.
 **************************************************************** {COPYRIGHT-END}
-->
<Jobs xmlns="http://www.ibm.com/TUAMJobs.xsd">
    <Job    id="AIXAA"
            description="Daily collection"
            active="true"
            joblogWriteToDB="false"
            joblogWriteToTextFile="true"
            joblogWriteToXMLFile="true"
            joblogShowStepOutput="true"
            joblogShowStepParameters="true"
            processPriorityClass="Low"
            smtpServer="mail.ITUAMCustomerCompany.com"
            smtpFrom="ITUAM@ITUAMCustomerCompany.com"
            smtpTo="John.ITUAMUser@ITUAMCustomerCompany.com"
            stopOnProcessFailure="false">
        <Process id="AIXAA"
                    description="Process for AIXAA data collection"
                    joblogShowStepOutput="true"
                    joblogShowStepParameters="true"
                    active="true">
      <Steps stopOnStepFailure="true">
        <Step id="Integrator"
            type="ConvertToCSR"
                        programName="integrator"
                        programType="java"
                        active="true">
        <Integrator>
          <Input name="AIXAAInput" active="true">
            <Files>
              <File name="%CollectorLogs%/AACCT_1/sample/aacct1_%LogDate_End%.txt" />
              <File name="%CollectorLogs%/AACCT_4/sample/aacct4_%LogDate_End%.txt" />
              <File name="%CollectorLogs%/AACCT_6/sample/aacct6_%LogDate_End%.txt" />
              <File name="%CollectorLogs%/AACCT_7/sample/aacct7_%LogDate_End%.txt" />
              <File name="%CollectorLogs%/AACCT_8/sample/aacct8_%LogDate_End%.txt" />
<!--              <File name="%CollectorLogs%/AACCT_10/sample/aacct10_%LogDate_End%.txt" /> -->
```

```
              <File name="%CollectorLogs%/AACCT_11/sample/aacct11_%LogDate_End%.txt" />
<!--              <File name="%CollectorLogs%/AACCT_14/sample/aacct14_%LogDate_End%.txt" /> -->
              <File name="%ProcessFolder%/exception.txt" type="exception" />
            </Files>
          </Input>

          <Stage name="CreateIdentifierFromTable" active="true">
            <Identifiers>
              <Identifier name="Account_Code_TMP">
                <FromIdentifiers>
                  <FromIdentifier name="SYSTEM_ID" offset="1" length="12" />
                </FromIdentifiers>
              </Identifier>
            </Identifiers>
            <Files>
              <File name="/opt/ibm/tuam/processes/accttabl.txt" type="table" />
              <File name="Exception.txt" type="exception" format="CSROutput" />
            </Files>
            <Parameters>
              <Parameter exceptionProcess="true" />
              <Parameter sort="true" />
              <Parameter upperCase="false" />
              <Parameter writeNoMatch="false" />
              <Parameter modifyIfExists="true" />
            </Parameters>
          </Stage>

          <Stage name="CreateIdentifierFromIdentifiers" active="true">
            <Identifiers>
              <Identifier name="Account_Code">
                <FromIdentifiers>
                  <FromIdentifier name="Account_Code_TMP" offset="1" length="28" />
                  <FromIdentifier name="SYSTEM_ID" offset="1" length="32" />
                </FromIdentifiers>
              </Identifier>
            </Identifiers>
            <Parameters>
              <Parameter keepLength="true" />
              <Parameter modifyIfIexists="true" />
            </Parameters>
          </Stage>

          <Stage name="DropFields" active="true">
            <Fields>
              <Field name="AAID0101"/>
              <Field name="AAID0102"/>
              <Field name="AAID0103"/>
              <Field name="AAID0403"/>
            </Fields>
```

```
         </Stage>

          <Stage name="CSRPlusOutput" active="true">
            <Files>
              <File name="%ProcessFolder%/server1/%LogDate_End%.txt" />
            </Files>
          </Stage>
        </Integrator>
</Step>

        <Step   id="Scan"
                description="Scan AIXAA"
                type="Process"
                programName="Scan"
                programType="java"
                active="true">
                <Parameters>
                   <Parameter retainFileDate="false"/>
                   <Parameter allowMissingFiles="false"/>
                   <Parameter allowEmptyFiles="false"/>
                   <Parameter useStepFiles="false"/>
                </Parameters>
        </Step>

        <Step   id="Process"
                        description="Standard Processing for AIXAA"
                        type="Process"
                        programName="Bill"
                        programType="java"
                        active="true">
            <Bill>
              <Parameters>
                       <Parameter inputFile="CurrentCSR.txt"/>
              </Parameters>
            </Bill>
        </Step>
        <Step   id="DatabaseLoad"
                description="Database Load for AIXAA"
                type="Process"
                programName="DBLoad"
                programType="java"
                active="true">
                <DBLoad>
                  <Parameters>
                  </Parameters>
                </DBLoad>
        </Step>
        <Step   id="Cleanup"
```

```
                description="Cleanup AIXAA"
                type="Process"
                programName="Cleanup"
                programType="java"
                active="false">
                <Parameters>
                    <Parameter DaysToRetainFiles="45"/>
                    <Parameter cleanSubfolders="true"/>
                </Parameters>
        </Step>
        </Steps>
        </Process>


    </Job>
</Jobs>
```

During the processing, we make a temporary copy of the template XML file for each AIX server, and change the job names and paths to data files to match that AIX server.

5.7.3 Script for collecting and loading the AIX accounting data

The script we use for collecting and loading the AIX accounting data is shown in Example 5-6. The script is intended to be run directly before midnight on the Tivoli Usage and Accounting Manager server.

Example 5-6 ituam_collect_data.sh

```
#!/bin/ksh
# This script is supposed to be started just before midnight on the TUAM server

DATETOGET=$(date +%Y%m%d)
TUAMHOSTS=/usr/local/bin/ituam_hosts
TUAMJOBS=/opt/ibm/tuam/jobfiles
JOBRUNNER=/opt/ibm/tuam/bin/startJobRunner.sh
CLIENTUSER=tuam
CLIENTPATH=/opt/ibm/tuam/collectors/unix/CS_input_source
SERVERPATH=/opt/ibm/tuam/logs/collectors
RECORDTYPES="1 4 6 7 8 11"

# Wait for the clock to roll over and the daily collection scripts
# on the clients to finish

sleep 600

# Get data from all the clients
```

```
cat $TUAMHOSTS|while read CLIENT
do
  echo Getting data from ${CLIENT}
  for TYPE in ${RECORDTYPES}
  do
    if [ ! -d ${SERVERPATH}/AACCT_${TYPE}/${CLIENT} ]
    then
      mkdir -p ${SERVERPATH}/AACCT_${TYPE}/${CLIENT}
    fi
    scp ${CLIENTUSER}@${CLIENT}:${CLIENTPATH}/aacct${TYPE}_${DATETOGET}.txt \
      ${SERVERPATH}/AACCT_${TYPE}/${CLIENT}/
  done
done

# Load data from all the clients

cat $TUAMHOSTS|while read CLIENT
do
  echo Loading data from ${CLIENT}
# First create a modified jobfile for the current host
  cat ${TUAMJOBS}/SampleAIXAA.xml|sed "s/sample/${CLIENT}/g" | \
    sed "s/\"AIXAA\"/\"AIXAA-${CLIENT}\"/g" > ${TUAMJOBS}/temp.xml
  ${JOBRUNNER} temp.xml
done
```

5.7.4 Monitoring the Tivoli Usage and Accounting Manager jobs

The Tivoli Usage and Accounting Manager jobs can be monitored from the Integrated Solutions Console, by selecting **Usage and Accounting Manager** → **Chargeback Maintenance** → **Job Runner** → **Log Files**, as shown in Figure 5-5 on page 253.

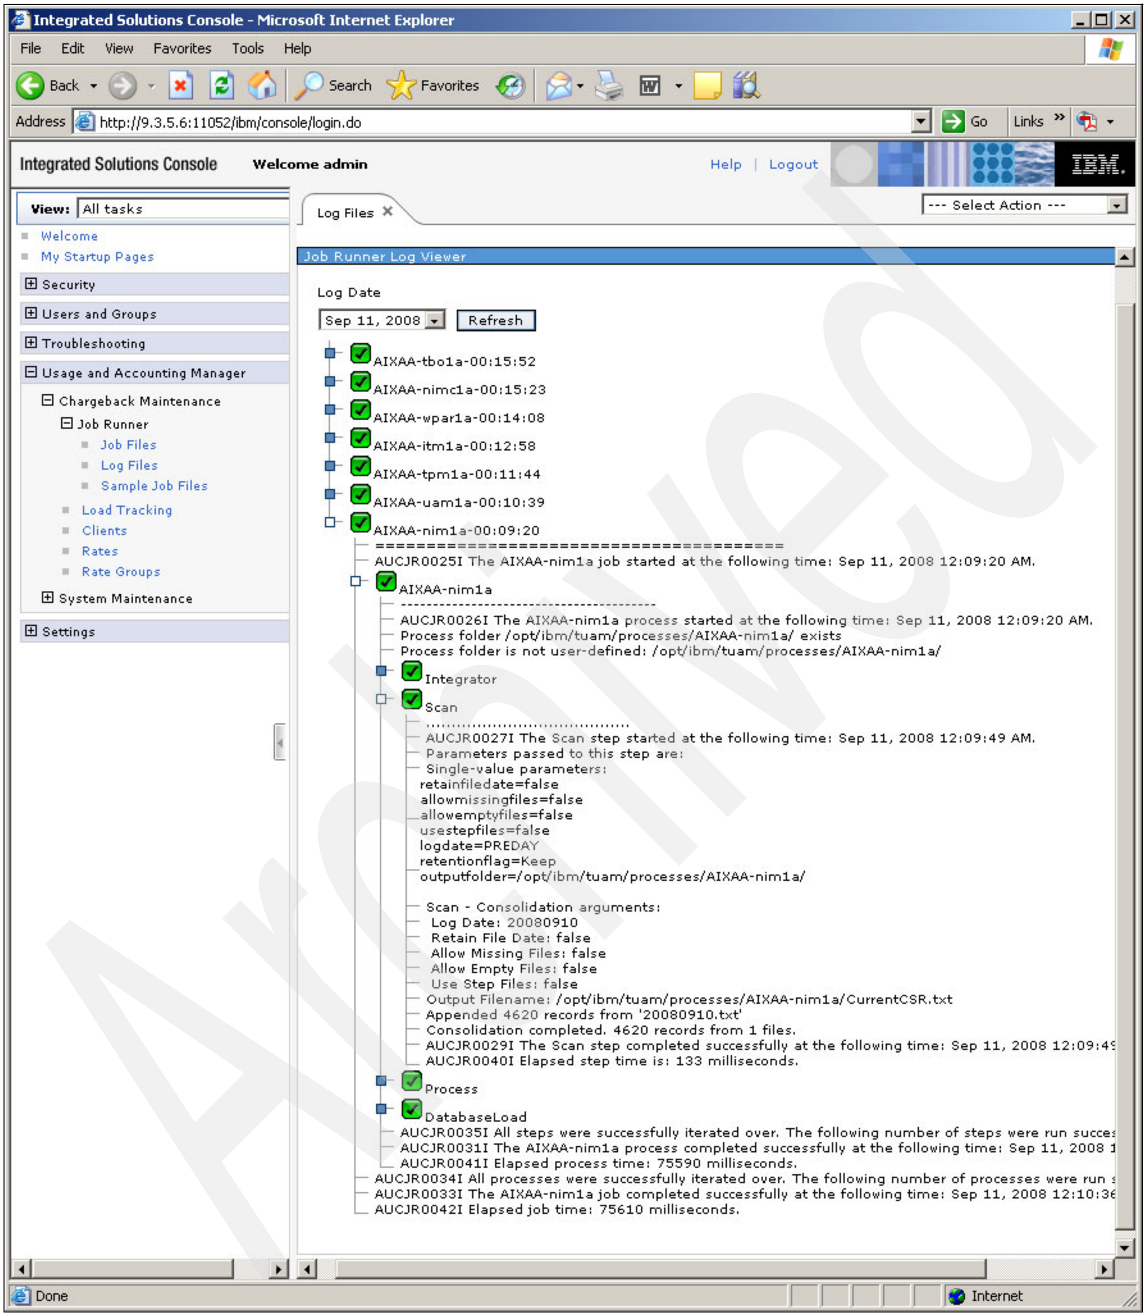

Log Date: Sep 11, 2008

```
AIXAA-tbo1a-00:15:52
AIXAA-nimc1a-00:15:23
AIXAA-wpar1a-00:14:08
AIXAA-itm1a-00:12:58
AIXAA-tpm1a-00:11:44
AIXAA-uam1a-00:10:39
AIXAA-nim1a-00:09:20
  ========================================
  AUCJR0025I The AIXAA-nim1a job started at the following time: Sep 11, 2008 12:09:20 AM.
    AIXAA-nim1a
    ----------------------------------------
    AUCJR0026I The AIXAA-nim1a process started at the following time: Sep 11, 2008 12:09:20 AM.
    Process folder /opt/ibm/tuam/processes/AIXAA-nim1a/ exists
    Process folder is not user-defined: /opt/ibm/tuam/processes/AIXAA-nim1a/
      Integrator
      Scan
      ........................................
      AUCJR0027I The Scan step started at the following time: Sep 11, 2008 12:09:49 AM.
      Parameters passed to this step are:
      Single-value parameters:
     retainfiledate=false
     allowmissingfiles=false
     allowemptyfiles=false
     usestepfiles=false
     logdate=PREDAY
     retentionflag=Keep
     outputfolder=/opt/ibm/tuam/processes/AIXAA-nim1a/

      Scan - Consolidation arguments:
       Log Date: 20080910
       Retain File Date: false
       Allow Missing Files: false
       Allow Empty Files: false
       Use Step Files: false
      Output Filename: /opt/ibm/tuam/processes/AIXAA-nim1a/CurrentCSR.txt
      Appended 4620 records from '20080910.txt'
      Consolidation completed. 4620 records from 1 files.
      AUCJR0029I The Scan step completed successfully at the following time: Sep 11, 2008 12:09:4
      AUCJR0040I Elapsed step time is: 133 milliseconds.
      Process
      DatabaseLoad
    AUCJR0035I All steps were successfully iterated over. The following number of steps were run succes
    AUCJR0031I The AIXAA-nim1a process completed successfully at the following time: Sep 11, 2008 1
    AUCJR0041I Elapsed process time: 75590 milliseconds.
  AUCJR0034I All processes were successfully iterated over. The following number of processes were run
  AUCJR0033I The AIXAA-nim1a job completed successfully at the following time: Sep 11, 2008 12:10:3
  AUCJR0042I Elapsed job time: 75610 milliseconds.
```

Figure 5-5 Tivoli Usage and Accounting Manager log files

5.8 Cost reporting and chargeback

In the Tivoli Usage and Accounting Manager, the cost reporting and chargeback functions are provided by the report server component. The report server reads the collected data from the Tivoli Usage and Accounting Manager database.

5.8.1 Reporting server overview

When you open a Web browser on the Tivoli Usage and Accounting Manager reporting server, you see the initial window as shown in Figure 5-6.

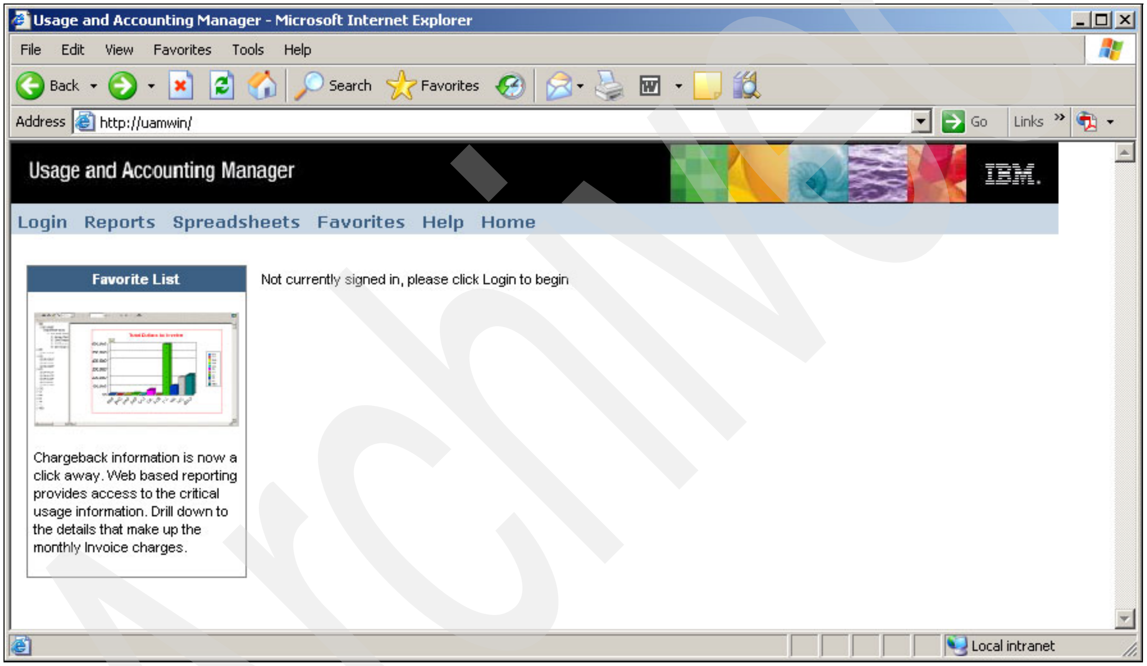
Usage and Accounting Manager - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Address http://uamwin/
Go Links
Usage and Accounting Manager
IBM.
Login Reports Spreadsheets Favorites Help Home
Favorite List
Local intranet

Not currently signed in, please click Login to begin

Chargeback information is now a click away. Web based reporting provides access to the critical usage information. Drill down to the details that make up the monthly Invoice charges.

Figure 5-6 Tivoli Usage and Accounting Manager reporting server initial window

To be able to run reports, you have to log in to the reporting server by choosing **Login** from the menu. The Tivoli Usage and Accounting Manager login window is shown in Figure 5-7 on page 255.

After installation, the system contains only the default user ID, admin, which has access to all the data and reports. You can create more user IDs, as required. Creating new user IDs is described in 5.9, “Providing resource utilization view to tenants” on page 266.

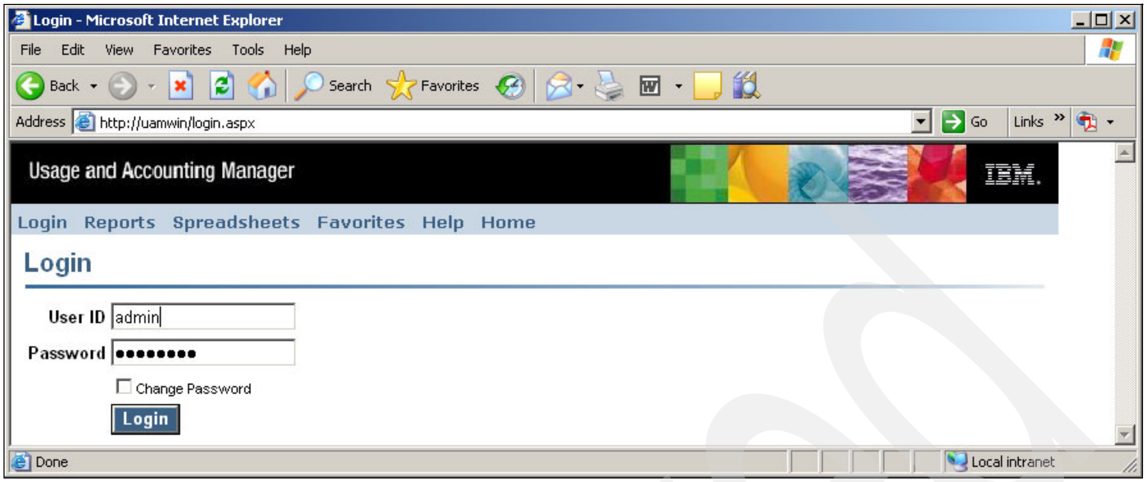
Login - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Address http://uamwin/login.aspx
Go Links
Usage and Accounting Manager
IBM.
Login Reports Spreadsheets Favorites Help Home
Done
Local intranet

Login

User ID admin

Password ●●●●●●●●●

□ Change Password

Login

Figure 5-7 Tivoli Usage and Accounting Manager login window

After you have logged in to the reporting server, you can run either reports or spreadsheets. In addition, you can choose the reports and spreadsheets you want to include in your Favorites menu. The options are described in this section.

Reports

The reports function allows you to create reports from the Tivoli Usage and Accounting Manager data. The reports can be later exported to a file in either Adobe® Acrobat® or Microsoft Excel formats. You can find a list of predefined reports by selecting **Reports** → **Run Reports**. A partial list of reports is shown in Figure 5-8 on page 256. You can also define your own report format by selecting **Reports** → **Create a Report**. The reports can be based on resource usage, cost, or both. Certain reports, such as CICS® or SQL Server reports, are not relevant in AIX environments.

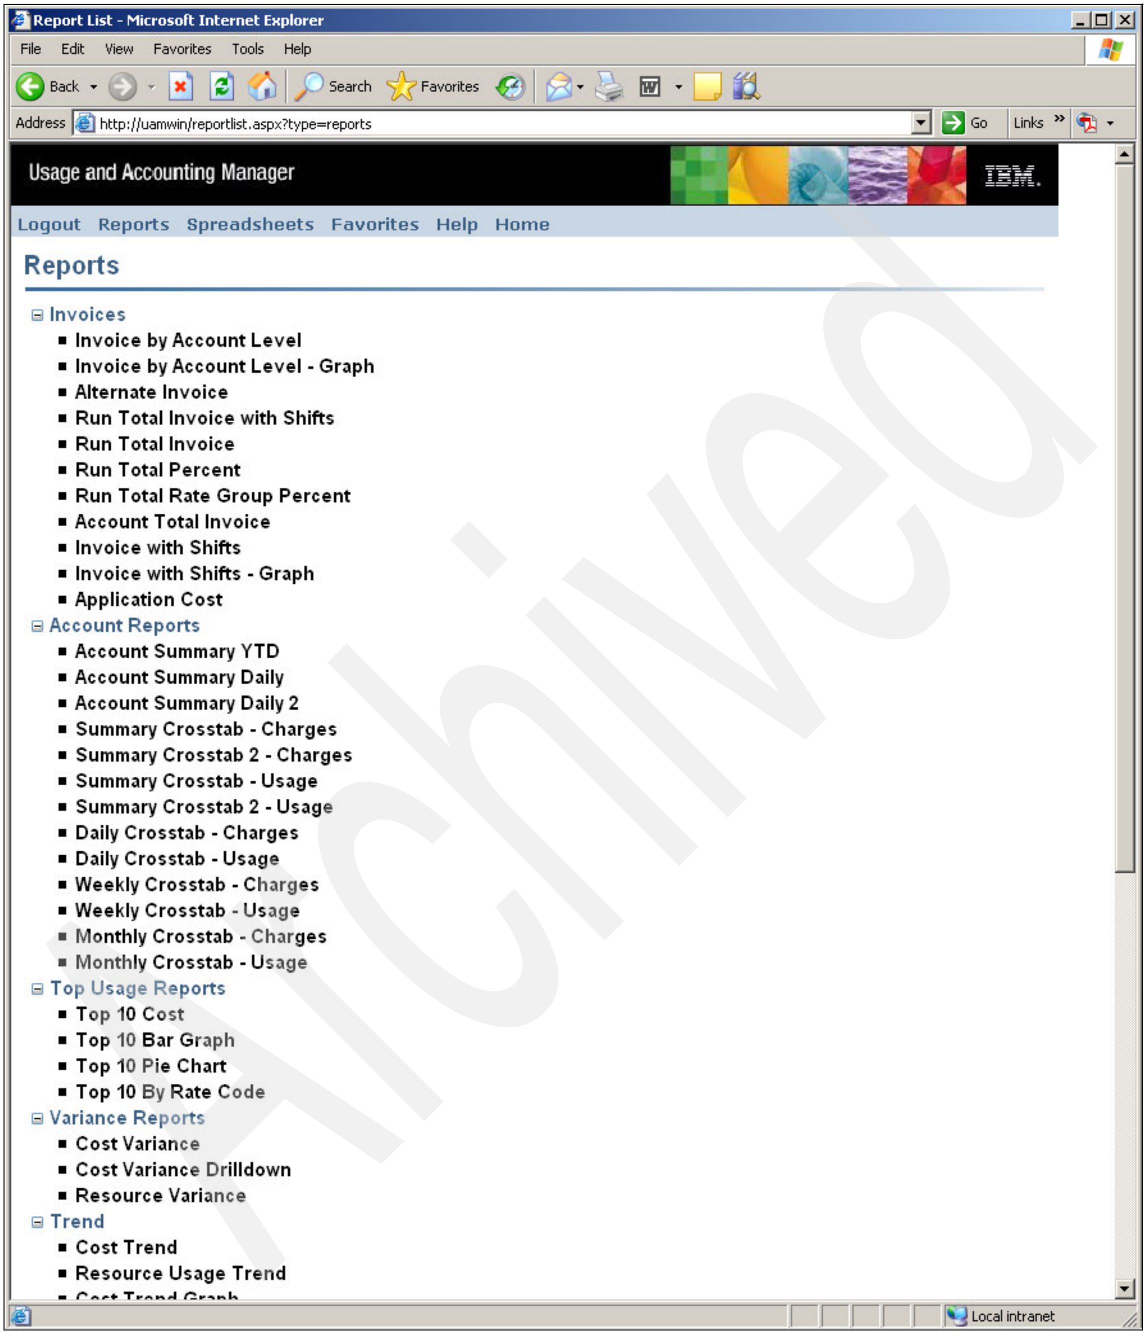
Report List - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back
Search
Favorites
Address http://uamwin/reportlist.aspx?type=reports
Go
Links
Usage and Accounting Manager
IBM.
Logout Reports Spreadsheets Favorites Help Home
Local intranet

Reports

Invoices

- Invoice by Account Level
- Invoice by Account Level - Graph
- Alternate Invoice
- Run Total Invoice with Shifts
- Run Total Invoice
- Run Total Percent
- Run Total Rate Group Percent
- Account Total Invoice
- Invoice with Shifts
- Invoice with Shifts - Graph
- Application Cost

Account Reports

- Account Summary YTD
- Account Summary Daily
- Account Summary Daily 2
- Summary Crosstab - Charges
- Summary Crosstab 2 - Charges
- Summary Crosstab - Usage
- Summary Crosstab 2 - Usage
- Daily Crosstab - Charges
- Daily Crosstab - Usage
- Weekly Crosstab - Charges
- Weekly Crosstab - Usage
- Monthly Crosstab - Charges
- Monthly Crosstab - Usage

Top Usage Reports

- Top 10 Cost
- Top 10 Bar Graph
- Top 10 Pie Chart
- Top 10 By Rate Code

Variance Reports

- Cost Variance
- Cost Variance Drilldown
- Resource Variance

Trend

- Cost Trend
- Resource Usage Trend
- Cost Trend Graph

Figure 5-8 Partial list of predefined Tivoli Usage and Accounting Manager reports

Spreadsheets

The spreadsheets function allows you to export data from the Tivoli Usage and Accounting Manager database into an Microsoft Excel file. You can find a list of predefined formats by selecting **Spreadsheets** → **Run Spreadsheets**. See Figure 5-9. You can also define your own spreadsheet format by selecting **Spreadsheets** → **Create a Spreadsheet**. The spreadsheets can be based on resource usage, cost, or both.

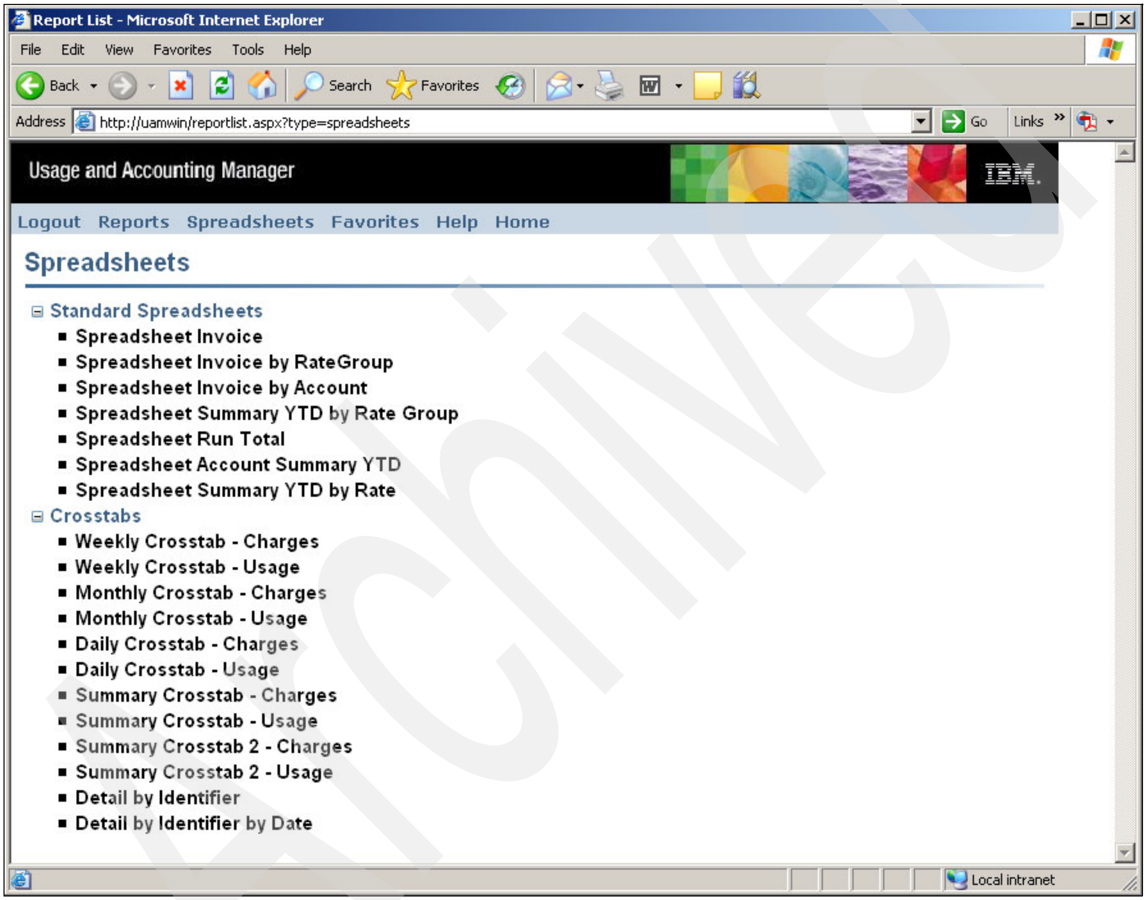
Report List - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Address http://uamwin/reportlist.aspx?type=spreadsheets
Go Links
Usage and Accounting Manager
IBM.
Logout Reports Spreadsheets Favorites Help Home
Local intranet

Spreadsheets

⊟ Standard Spreadsheets

- Spreadsheet Invoice
- Spreadsheet Invoice by RateGroup
- Spreadsheet Invoice by Account
- Spreadsheet Summary YTD by Rate Group
- Spreadsheet Run Total
- Spreadsheet Account Summary YTD
- Spreadsheet Summary YTD by Rate

⊟ Crosstabs

- Weekly Crosstab - Charges
- Weekly Crosstab - Usage
- Monthly Crosstab - Charges
- Monthly Crosstab - Usage
- Daily Crosstab - Charges
- Daily Crosstab - Usage
- Summary Crosstab - Charges
- Summary Crosstab - Usage
- Summary Crosstab 2 - Charges
- Summary Crosstab 2 - Usage
- Detail by Identifier
- Detail by Identifier by Date

Figure 5-9 List of predefined Tivoli Usage and Accounting Manager spreadsheets

Favorites

The favorites function allows you to select the reports or spreadsheets you want to have in your Favorite Reports and Favorite Spreadsheets menus. These menus enable you to more easily find the reports or spreadsheets you use most.

You can select your favorite reports by selecting **Favorites** → **Reports**, or select your favorite spreadsheets by selecting **Favorites** → **Spreadsheets**, and then clicking **Save**, as shown in Figure 5-10.

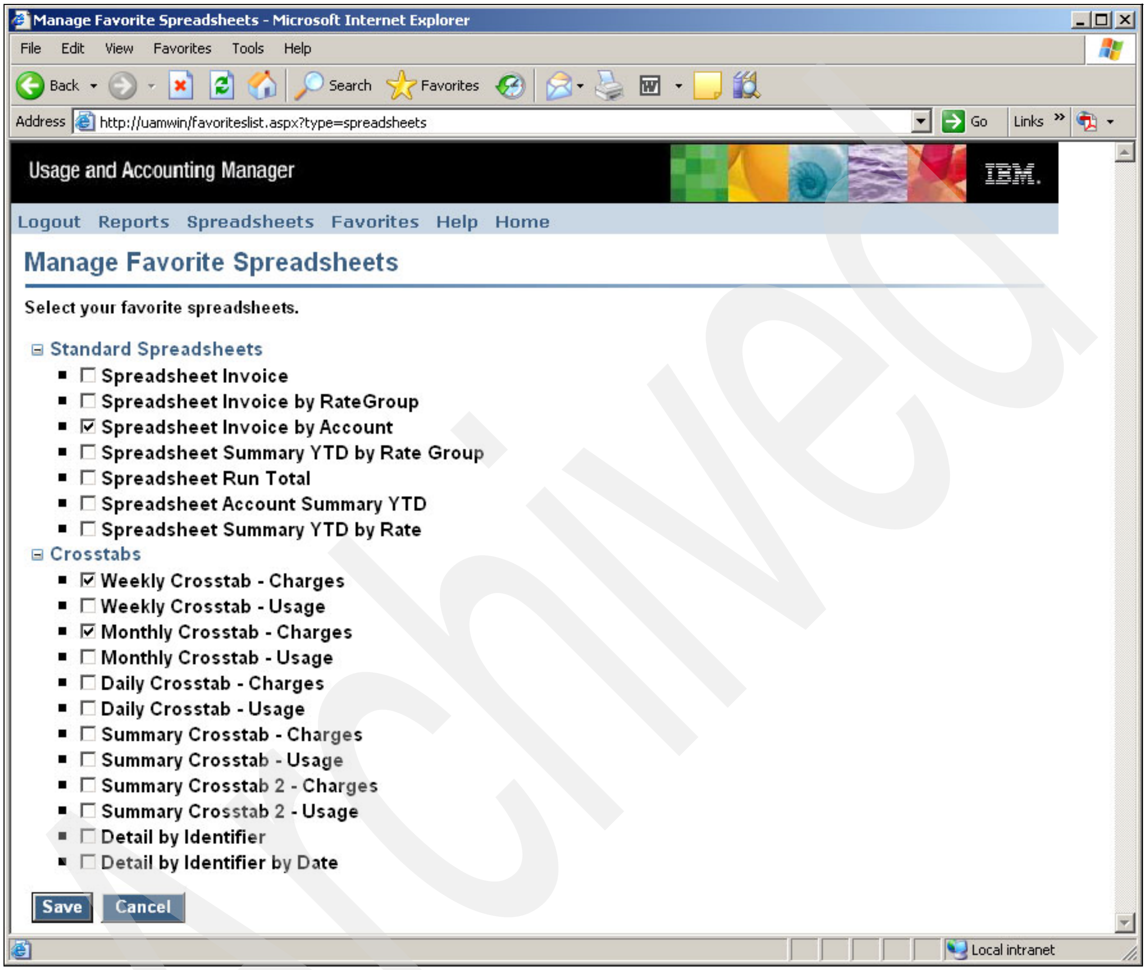
Manage Favorite Spreadsheets - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Address http://uamwin/favoriteslist.aspx?type=spreadsheets
Go Links
Usage and Accounting Manager
IBM.
Logout Reports Spreadsheets Favorites Help Home
Local intranet

Manage Favorite Spreadsheets

Select your favorite spreadsheets.

⊟ Standard Spreadsheets

- ☐ Spreadsheet Invoice
- ☐ Spreadsheet Invoice by RateGroup
- ☑ Spreadsheet Invoice by Account
- ☐ Spreadsheet Summary YTD by Rate Group
- ☐ Spreadsheet Run Total
- ☐ Spreadsheet Account Summary YTD
- ☐ Spreadsheet Summary YTD by Rate

⊟ Crosstabs

- ☑ Weekly Crosstab - Charges
- ☐ Weekly Crosstab - Usage
- ☑ Monthly Crosstab - Charges
- ☐ Monthly Crosstab - Usage
- ☐ Daily Crosstab - Charges
- ☐ Daily Crosstab - Usage
- ☐ Summary Crosstab - Charges
- ☐ Summary Crosstab - Usage
- ☐ Summary Crosstab 2 - Charges
- ☐ Summary Crosstab 2 - Usage
- ☐ Detail by Identifier
- ☐ Detail by Identifier by Date

Save Cancel

Figure 5-10 Choosing the favorite spreadsheets in Tivoli Usage and Accounting Manager

After defining the favorite reports, you can select a report from the list by selecting **Reports** → **Favorite Reports**. After defining the favorite spreadsheets, you can select a spreadsheet from the list of favorite spreadsheets by selecting **Spreadsheets** → **Favorite Spreadsheets**. See Figure 5-11 on page 259.

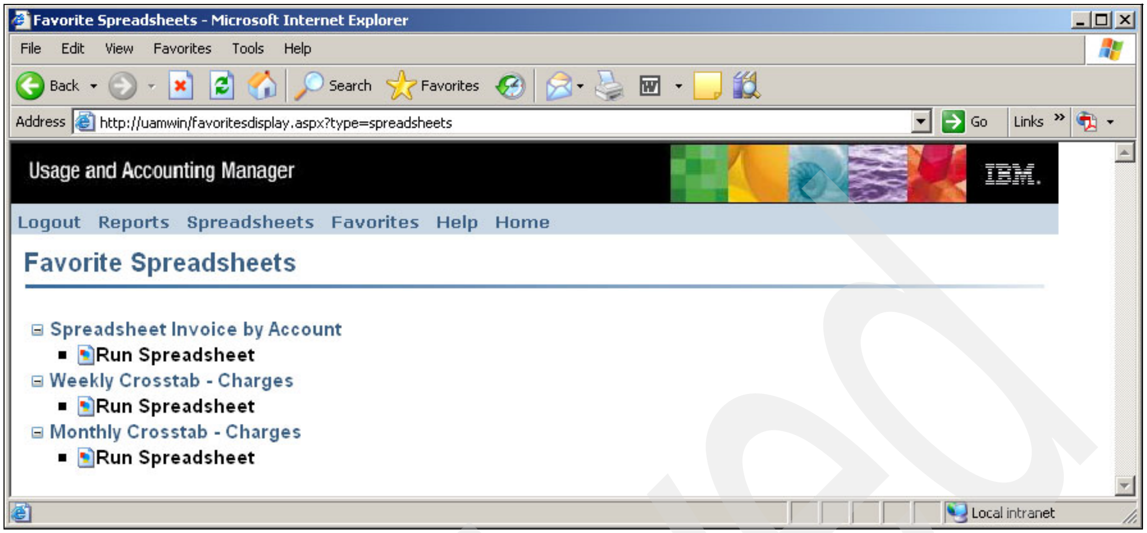
Favorite Spreadsheets - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Address http://uamwin/favoritesdisplay.aspx?type=spreadsheets
Go Links
Usage and Accounting Manager
IBM.
Logout Reports Spreadsheets Favorites Help Home
Local intranet

Favorite Spreadsheets

- Spreadsheet Invoice by Account
  - Run Spreadsheet
- Weekly Crosstab - Charges
  - Run Spreadsheet
- Monthly Crosstab - Charges
  - Run Spreadsheet

Figure 5-11 Choosing from the favorite Spreadsheets

5.8.2 Creating reports and spreadsheets

After selecting a report or spreadsheet from the list, you may choose parameters, which usually include:

- Account Code Level, based on the account code structure created in 5.5.1, “Defining the account code structure” on page 237, and shown in Figure 5-12 on page 260.
- Starting Account Code and Ending Account Code, based on the defined clients as shown in Figure 5-13 on page 260; or enter the range manually as Custom value
- Date Range, based on either the predefined values as shown in Figure 5-14 on page 261, or enter it manually as Custom value

The parameters available vary based on the report selected. For example, all invoices include the Invoice Number of the first invoice printed as an additional input field.

The spreadsheets are created in the same way as the reports, however, instead of viewing the report in a Web browser, you can save the spreadsheet as a file.

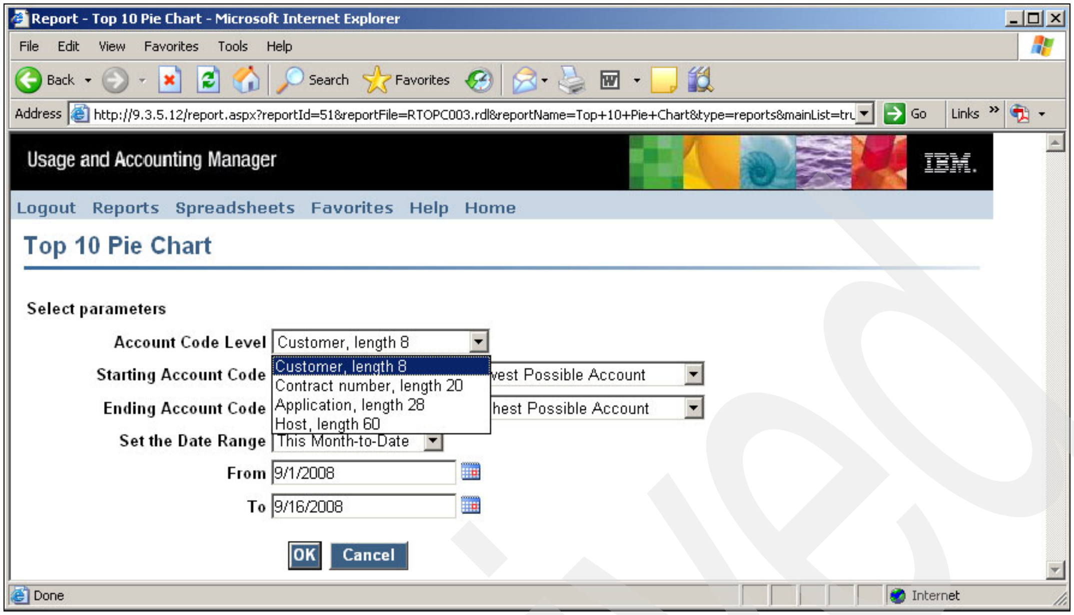
Report - Top 10 Pie Chart - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Go Links
Done
Internet
Usage and Accounting Manager
IBM.
Logout Reports Spreadsheets Favorites Help Home
Customer, length 8
Contract number, length 20
Application, length 28
Host, length 60
OK
Cancel

Address http://9.3.5.12/report.aspx?reportId=51&reportFile=RTOPC003.rdl&reportName=Top+10+Pie+Chart&type=reports&mainList=tru

Top 10 Pie Chart

Select parameters

Account Code Level: Customer, length 8

Starting Account Code: ...west Possible Account

Ending Account Code: ...hest Possible Account

Set the Date Range: This Month-to-Date

From: 9/1/2008

To: 9/16/2008

Figure 5-12 Choosing the Account Code Level

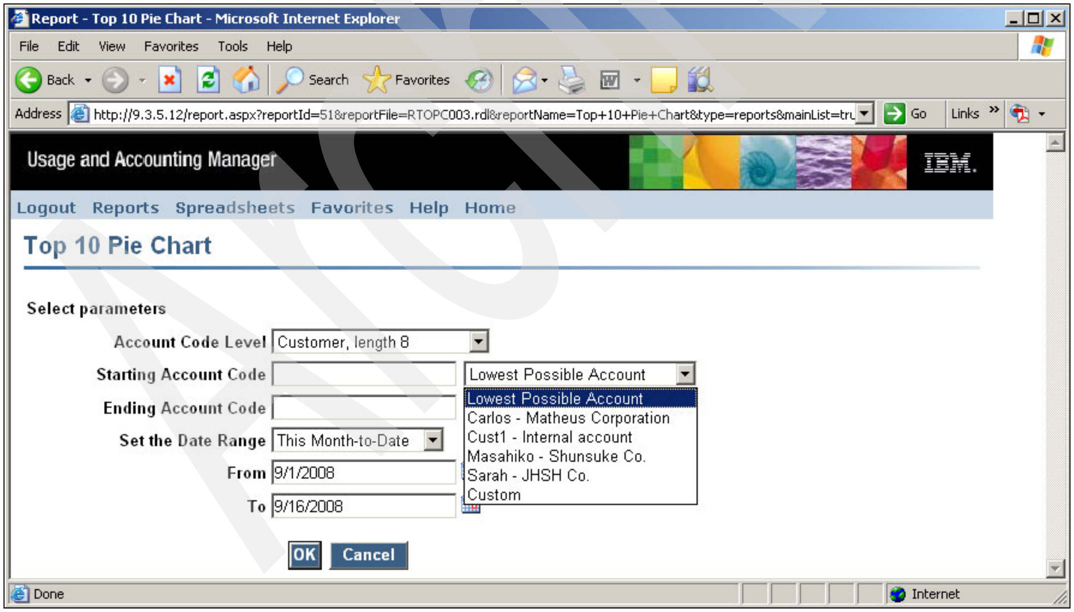
Report - Top 10 Pie Chart - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Go Links
Done
Internet
Usage and Accounting Manager
IBM.
Logout Reports Spreadsheets Favorites Help Home
Lowest Possible Account
Carlos - Matheus Corporation
Cust1 - Internal account
Masahiko - Shunsuke Co.
Sarah - JHSH Co.
Custom
OK
Cancel

Address http://9.3.5.12/report.aspx?reportId=51&reportFile=RTOPC003.rdl&reportName=Top+10+Pie+Chart&type=reports&mainList=tru

Top 10 Pie Chart

Select parameters

Account Code Level: Customer, length 8

Starting Account Code: Lowest Possible Account

Ending Account Code:

Set the Date Range: This Month-to-Date

From: 9/1/2008

To: 9/16/2008

Figure 5-13 Choosing the Starting Account Code

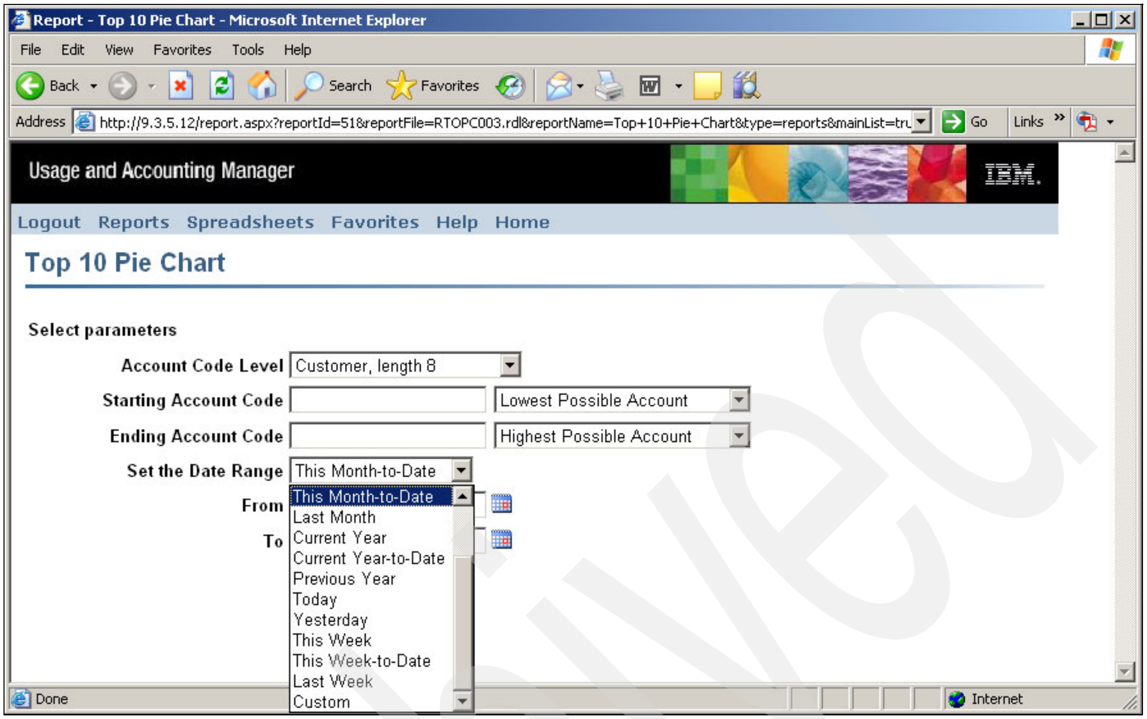

Report - Top 10 Pie Chart - Microsoft Internet Explorer

Address: http://9.3.5.12/report.aspx?reportId=51&reportFile=RTOPC003.rdl&reportName=Top+10+Pie+Chart&type=reports&mainList=tru

Usage and Accounting Manager

Top 10 Pie Chart

Select parameters

Account Code Level: Customer, length 8

Starting Account Code: Lowest Possible Account

Ending Account Code: Highest Possible Account

Set the Date Range: This Month-to-Date

Figure 5-14 Setting the date range

5.8.3 Accounting

Based on the data the reports use, the two basic types of reports in Tivoli Usage and Accounting Manager are:

- Reports based on resource usage metrics, such as the used CPU time
- Reports based on charges

Both types can be used for accounting purposes. Identifying the type of data that the report contains is usually easy to do from the name of the report.

Although the Crosstabs and especially the Trend reports are useful for capacity planning, they require at least several months of collected data to provide meaningful results.

As an example of the Tivoli Usage and Accounting Manager reporting, we show the Customer level Top 10 Pie Chart in Figure 5-15 on page 262.

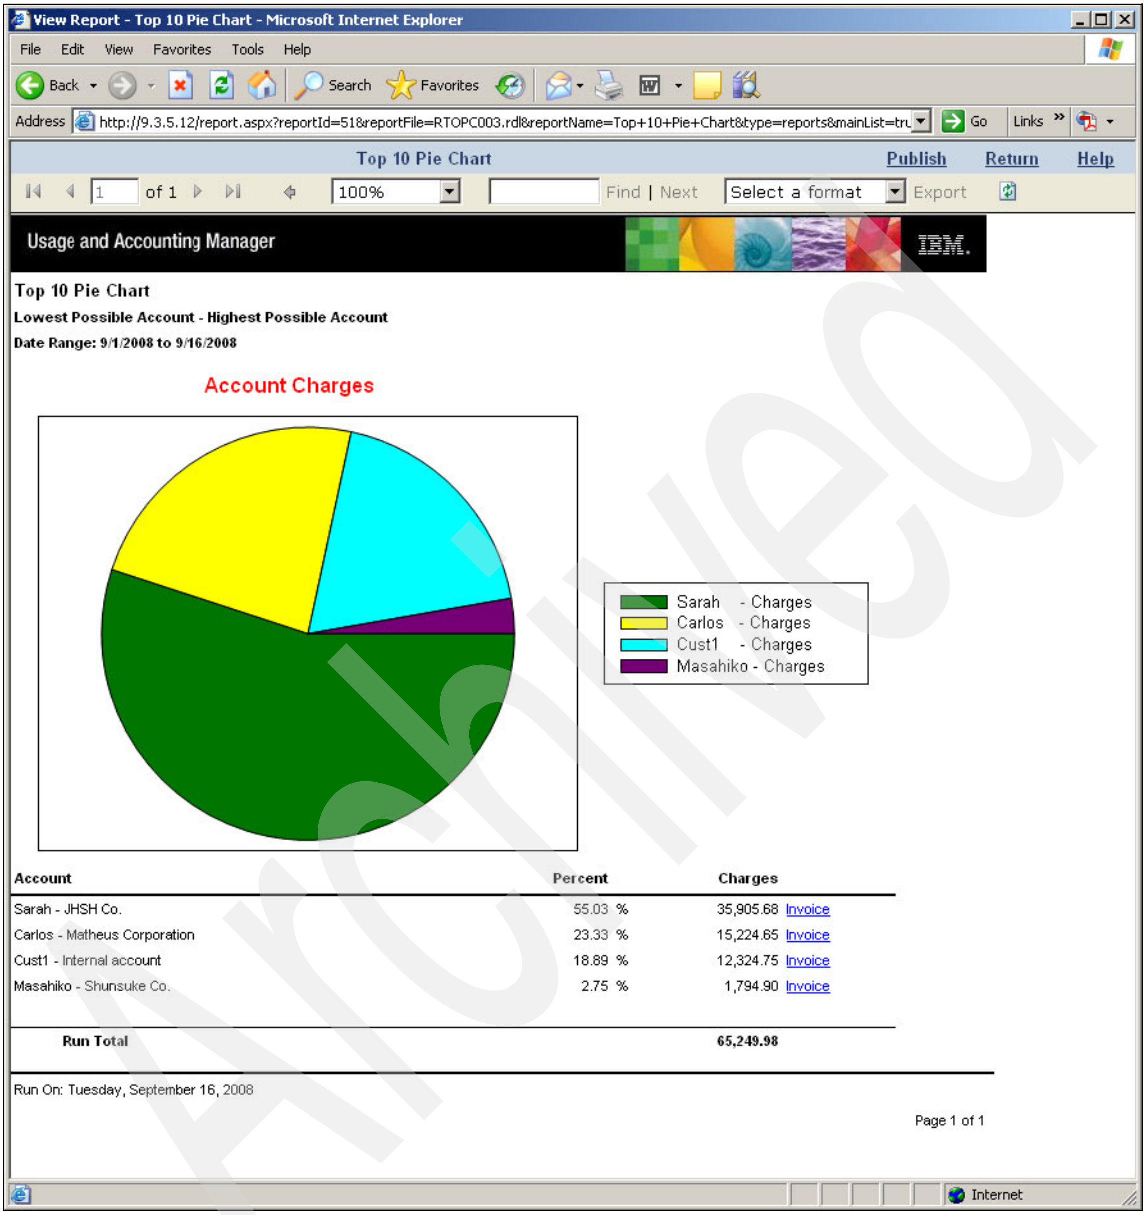

Address: http://9.3.5.12/report.aspx?reportId=51&reportFile=RTOPC003.rdl&reportName=Top+10+Pie+Chart&type=reports&mainList=tru

Top 10 Pie Chart

Usage and Accounting Manager

Top 10 Pie Chart

Lowest Possible Account - Highest Possible Account

Date Range: 9/1/2008 to 9/16/2008

Account Charges

🟩 Sarah - Charges
🟨 Carlos - Charges
🟦 Cust1 - Charges
🟪 Masahiko - Charges

Account	Percent	Charges	
Sarah - JHSH Co.	55.03 %	35,905.68	Invoice
Carlos - Matheus Corporation	23.33 %	15,224.65	Invoice
Cust1 - Internal account	18.89 %	12,324.75	Invoice
Masahiko - Shunsuke Co.	2.75 %	1,794.90	Invoice
Run Total		65,249.98	

Run On: Tuesday, September 16, 2008

Page 1 of 1

Figure 5-15 Tivoli Usage and Accounting Manager Top 10 Pie Chart

5.8.4 Chargeback and billing

The Invoice by Account Level is probably one of the most useful reports for chargeback purposes. You may choose the level of detail to include by your account code structure. The level can range from a single invoice for each tenant to having a separate invoice for each server.

We show an example invoice in Figure 5-16 on page 264 (part 1) and the continuation in Figure 5-17 on page 265 (part 2).

The values in the rate column are defined in the rate tables, as described in 5.5.4, “Setting up the rate tables” on page 239.

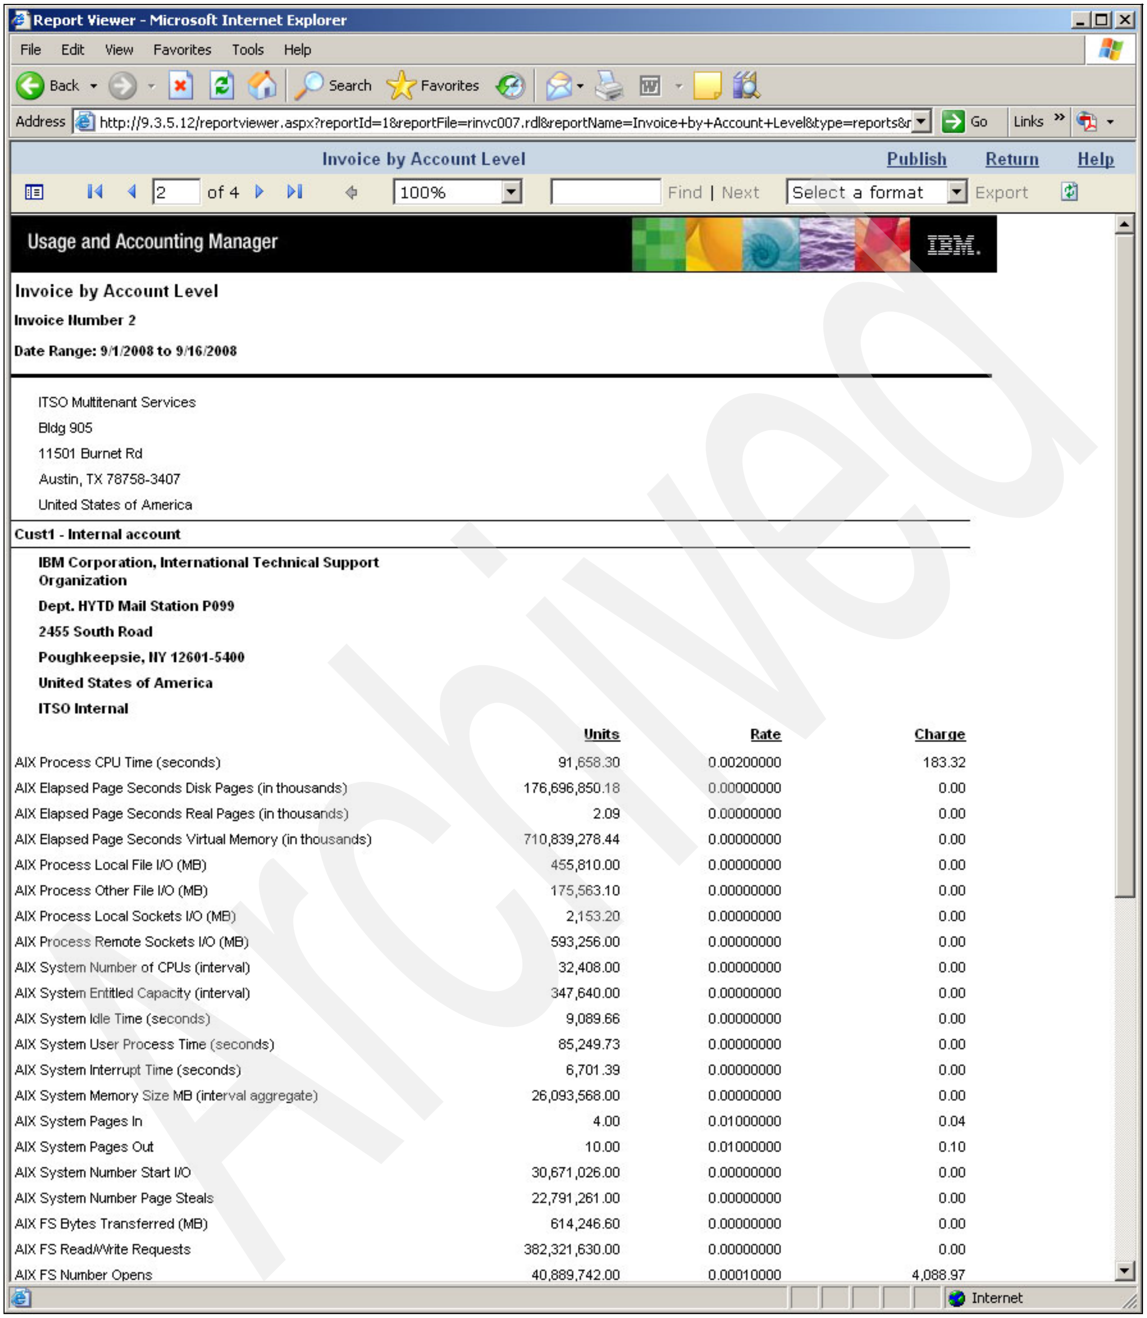

Report Viewer - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Address http://9.3.5.12/reportviewer.aspx?reportId=1&reportFile=rinvc007.rdl&reportName=Invoice+by+Account+Level&type=reports&r
Go Links
Invoice by Account Level
Publish Return Help
2 of 4
100%
Find | Next
Select a format
Export
Usage and Accounting Manager
IBM.
Internet

Invoice by Account Level

Invoice Number 2

Date Range: 9/1/2008 to 9/16/2008

ITSO Multitenant Services
Bldg 905
11501 Burnet Rd
Austin, TX 78758-3407
United States of America

Cust1 - Internal account

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400
United States of America
ITSO Internal

	Units	Rate	Charge
AIX Process CPU Time (seconds)	91,658.30	0.00200000	183.32
AIX Elapsed Page Seconds Disk Pages (in thousands)	176,696,850.18	0.00000000	0.00
AIX Elapsed Page Seconds Real Pages (in thousands)	2.09	0.00000000	0.00
AIX Elapsed Page Seconds Virtual Memory (in thousands)	710,839,278.44	0.00000000	0.00
AIX Process Local File I/O (MB)	455,810.00	0.00000000	0.00
AIX Process Other File I/O (MB)	175,563.10	0.00000000	0.00
AIX Process Local Sockets I/O (MB)	2,153.20	0.00000000	0.00
AIX Process Remote Sockets I/O (MB)	593,256.00	0.00000000	0.00
AIX System Number of CPUs (interval)	32,408.00	0.00000000	0.00
AIX System Entitled Capacity (interval)	347,640.00	0.00000000	0.00
AIX System Idle Time (seconds)	9,089.66	0.00000000	0.00
AIX System User Process Time (seconds)	85,249.73	0.00000000	0.00
AIX System Interrupt Time (seconds)	6,701.39	0.00000000	0.00
AIX System Memory Size MB (interval aggregate)	26,093,568.00	0.00000000	0.00
AIX System Pages In	4.00	0.01000000	0.04
AIX System Pages Out	10.00	0.01000000	0.10
AIX System Number Start I/O	30,671,026.00	0.00000000	0.00
AIX System Number Page Steals	22,791,261.00	0.00000000	0.00
AIX FS Bytes Transferred (MB)	614,246.60	0.00000000	0.00
AIX FS Read/Write Requests	382,321,630.00	0.00000000	0.00
AIX FS Number Opens	40,889,742.00	0.00010000	4,088.97

Figure 5-16 Tivoli Usage and Accounting Manager invoice (part 1)

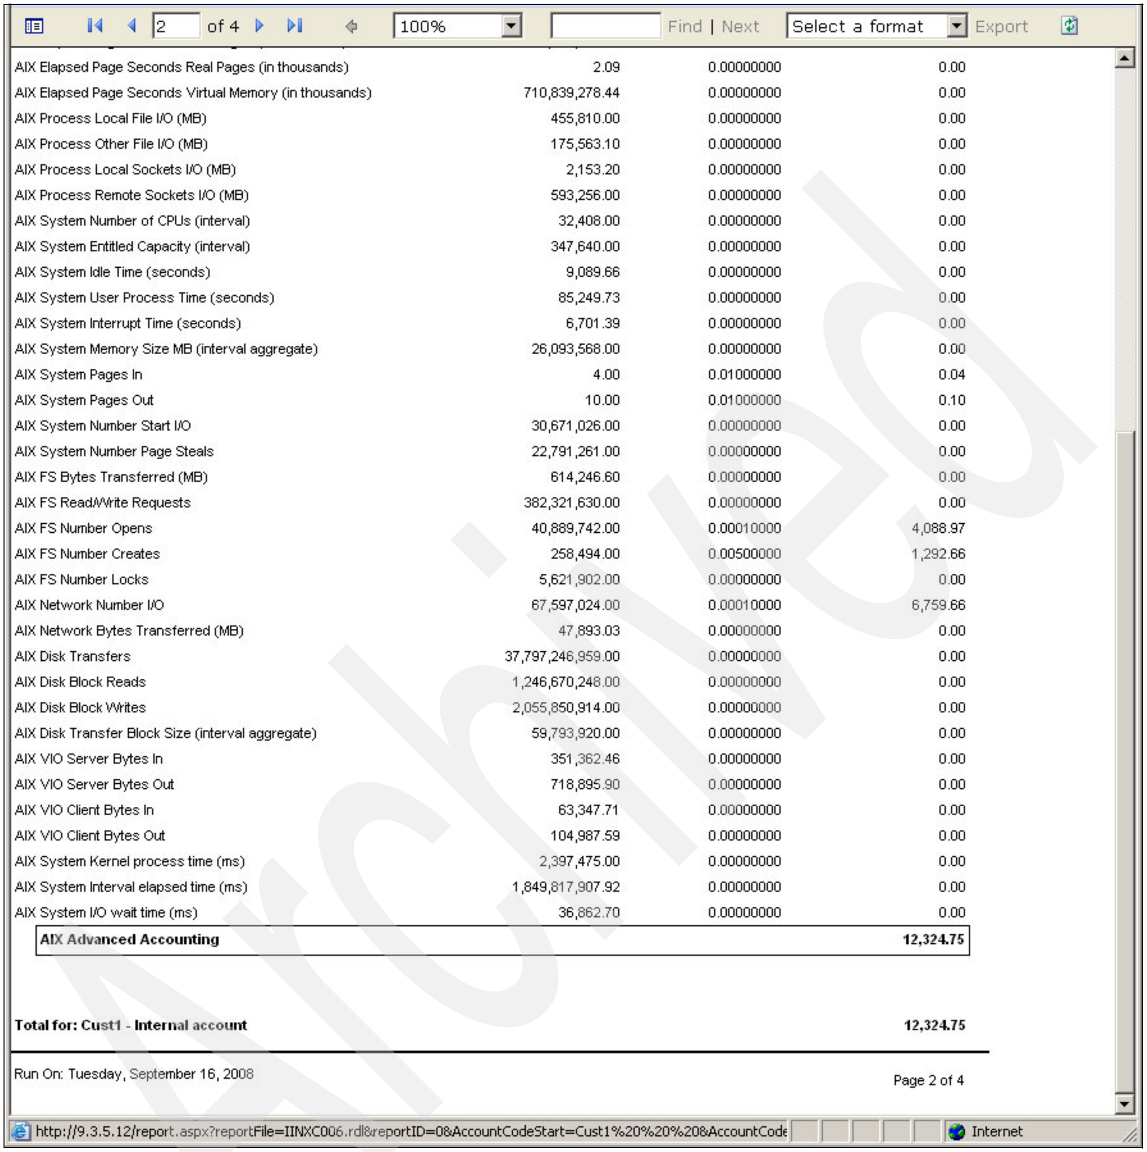

2
of 4
100%
Find | Next
Select a format
Export
http://9.3.5.12/report.aspx?reportFile=IINXC006.rdl&reportID=0&AccountCodeStart=Cust1%20%20%20&AccountCode
Internet

AIX Elapsed Page Seconds Real Pages (in thousands)	2.09	0.00000000	0.00
AIX Elapsed Page Seconds Virtual Memory (in thousands)	710,839,278.44	0.00000000	0.00
AIX Process Local File I/O (MB)	455,810.00	0.00000000	0.00
AIX Process Other File I/O (MB)	175,563.10	0.00000000	0.00
AIX Process Local Sockets I/O (MB)	2,153.20	0.00000000	0.00
AIX Process Remote Sockets I/O (MB)	593,256.00	0.00000000	0.00
AIX System Number of CPUs (interval)	32,408.00	0.00000000	0.00
AIX System Entitled Capacity (interval)	347,640.00	0.00000000	0.00
AIX System Idle Time (seconds)	9,089.66	0.00000000	0.00
AIX System User Process Time (seconds)	85,249.73	0.00000000	0.00
AIX System Interrupt Time (seconds)	6,701.39	0.00000000	0.00
AIX System Memory Size MB (interval aggregate)	26,093,568.00	0.00000000	0.00
AIX System Pages In	4.00	0.01000000	0.04
AIX System Pages Out	10.00	0.01000000	0.10
AIX System Number Start I/O	30,671,026.00	0.00000000	0.00
AIX System Number Page Steals	22,791,261.00	0.00000000	0.00
AIX FS Bytes Transferred (MB)	614,246.60	0.00000000	0.00
AIX FS Read/Write Requests	382,321,630.00	0.00000000	0.00
AIX FS Number Opens	40,889,742.00	0.00010000	4,088.97
AIX FS Number Creates	258,494.00	0.00500000	1,292.66
AIX FS Number Locks	5,621,902.00	0.00000000	0.00
AIX Network Number I/O	67,597,024.00	0.00010000	6,759.66
AIX Network Bytes Transferred (MB)	47,893.03	0.00000000	0.00
AIX Disk Transfers	37,797,246,959.00	0.00000000	0.00
AIX Disk Block Reads	1,246,670,248.00	0.00000000	0.00
AIX Disk Block Writes	2,055,850,914.00	0.00000000	0.00
AIX Disk Transfer Block Size (interval aggregate)	59,793,920.00	0.00000000	0.00
AIX VIO Server Bytes In	351,362.46	0.00000000	0.00
AIX VIO Server Bytes Out	718,895.90	0.00000000	0.00
AIX VIO Client Bytes In	63,347.71	0.00000000	0.00
AIX VIO Client Bytes Out	104,987.59	0.00000000	0.00
AIX System Kernel process time (ms)	2,397,475.00	0.00000000	0.00
AIX System Interval elapsed time (ms)	1,849,817,907.92	0.00000000	0.00
AIX System I/O wait time (ms)	36,862.70	0.00000000	0.00
AIX Advanced Accounting			**12,324.75**

Total for: Cust1 - Internal account **12,324.75**

Run On: Tuesday, September 16, 2008

Page 2 of 4

Figure 5-17 Tivoli Usage and Accounting Manager invoice (part 2)

5.9 Providing resource utilization view to tenants

In addition to using the Tivoli Usage and Accounting Manager for internal purposes, a service provider may choose to offer a limited resource utilization view to certain tenants. This way the tenants can track their resource utilization almost in real-time.

To use the service, the tenants have to be able to open an HTTP connection to a reporting server. This approach has obvious security ramifications, like any other Web service, which are outside the scope of this book.

5.9.1 Defining clients

When you provide information to the tenants, ensure that each tenant can only see the information concerning that tenant. In the Tivoli Usage and Accounting Manager, this kind of limited view is implemented by creating a client definition and giving the user group of the tenant access to only a certain client.

In our example, we create a client definition for account code Cust1, as shown in Figure 5-18 on page 267. This step means that this client definition matches any data for which the account code starts with Cust1. The tenant name is the first part of our account code, which limits access to servers of that one tenant. The clients can be defined in the Integrated Solutions Console by selecting **Usage and Accounting Manager** → **Chargeback Maintenance** → **Clients**.

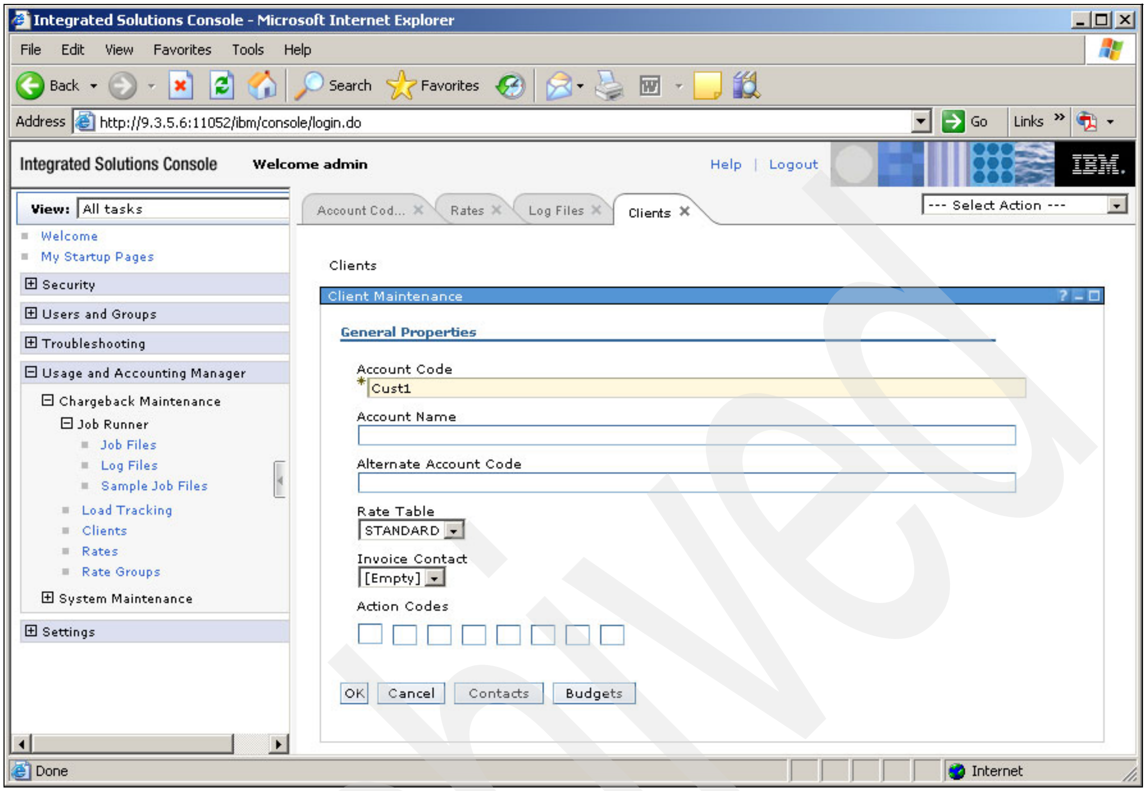

Account Code
* Cust1

Account Name

Alternate Account Code

Rate Table
STANDARD

Invoice Contact
[Empty]

Action Codes
□ □ □ □ □ □ □ □

Figure 5-18 Defining a client

5.9.2 Defining user groups

In the Tivoli Usage and Accounting Manager reporting server, access control is implemented on the user group level. This approach means that each user group has a list of clients it can access, and list of reports it can run.

In our example, we define a user group named Client1 to hold all users from tenant named Client1, as shown in Figure 5-19 on page 268. The user groups can be defined in the Integrated Solutions Console by selecting **Usage and Accounting Manager** → **System Maintenance** → **User Groups**.

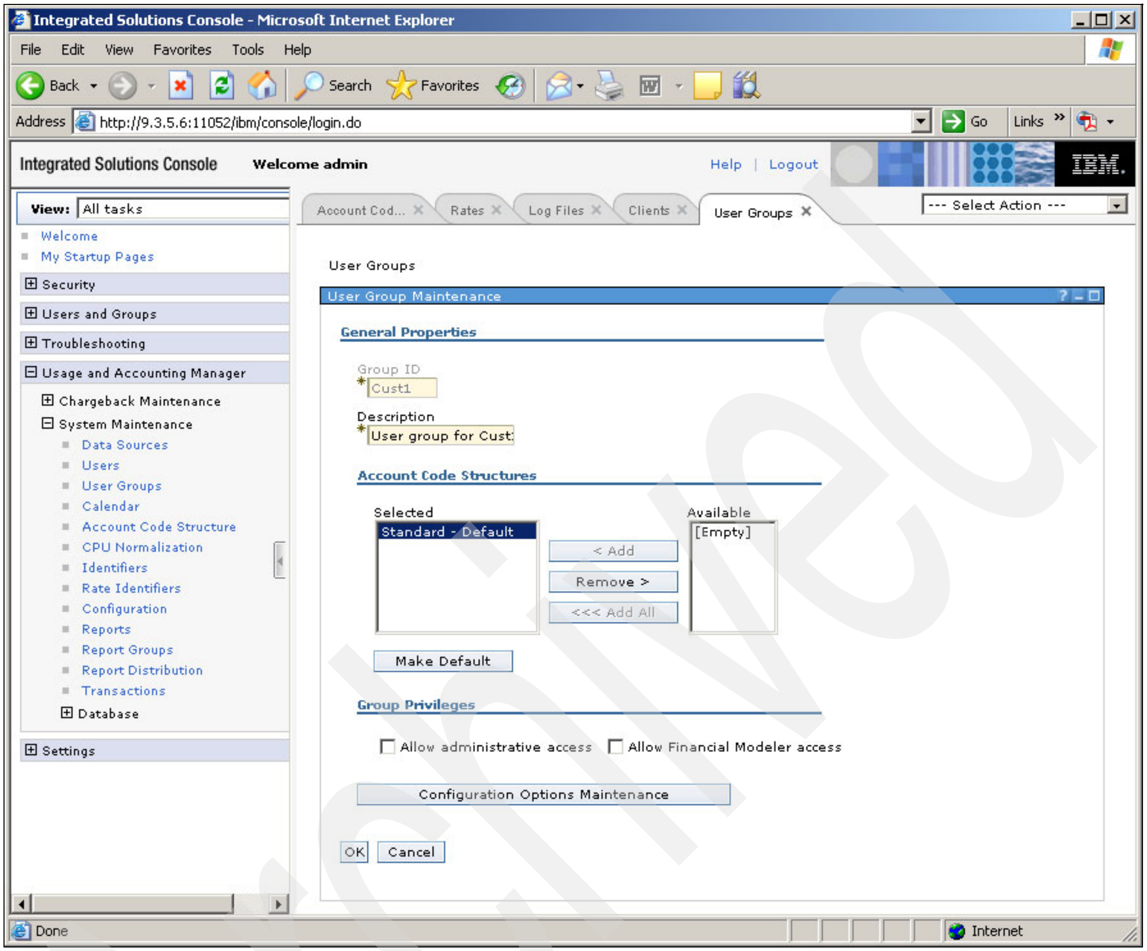

Group ID
* Cust1

Description
* User group for Cust:

Selected: Standard - Default

Available: [Empty]

□ Allow administrative access □ Allow Financial Modeler access

Figure 5-19 Defining a user group

By default, the new user group has access to all accounting data and all reports. We have to change the user group to only include the clients and reports we want the tenant to see, as shown in Figure 5-20 on page 269.

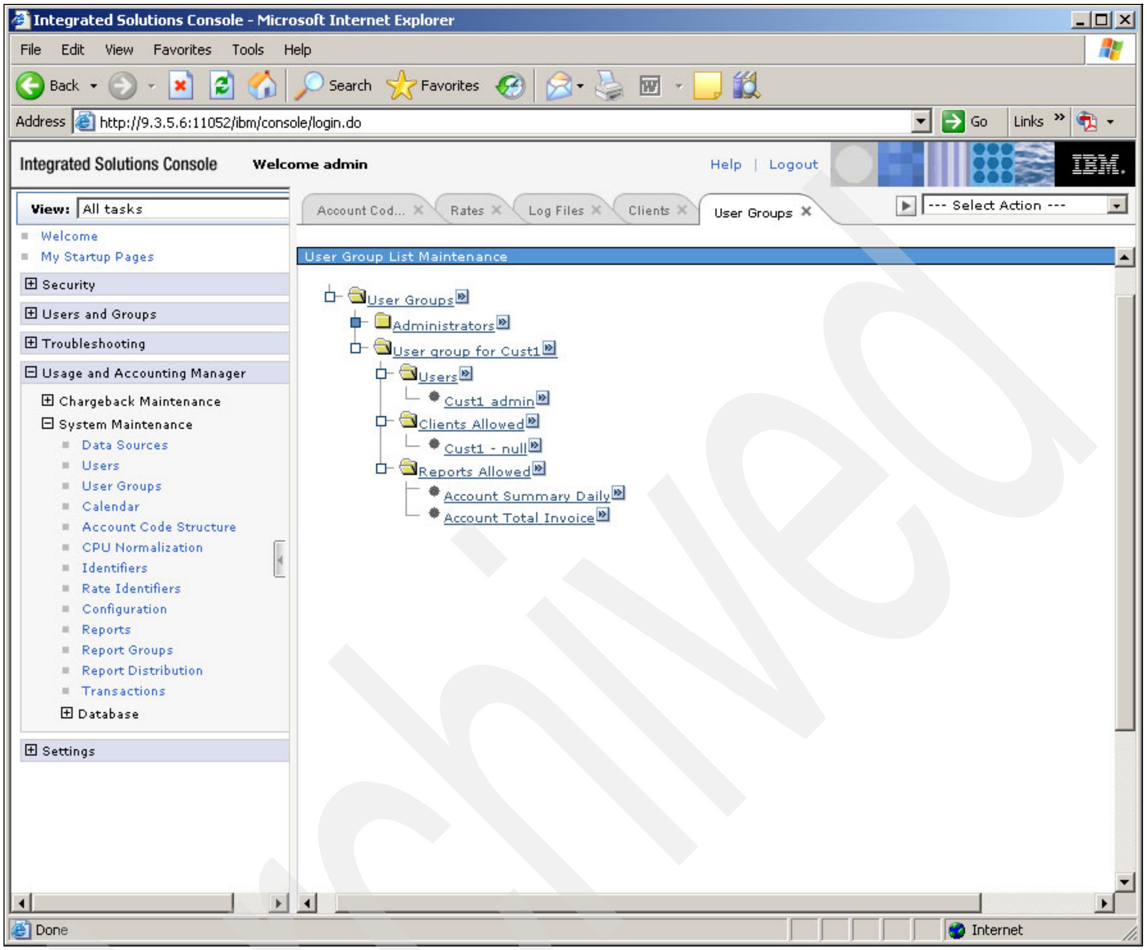

Figure 5-20 Changing user group privileges

5.9.3 Defining users

After we have defined the user group, we can define the user that will use these privileges, as shown in Figure 5-21. The users can be defined in the Integrated Solutions Console by selecting **Usage and Accounting Manager → System Maintenance → Users**.

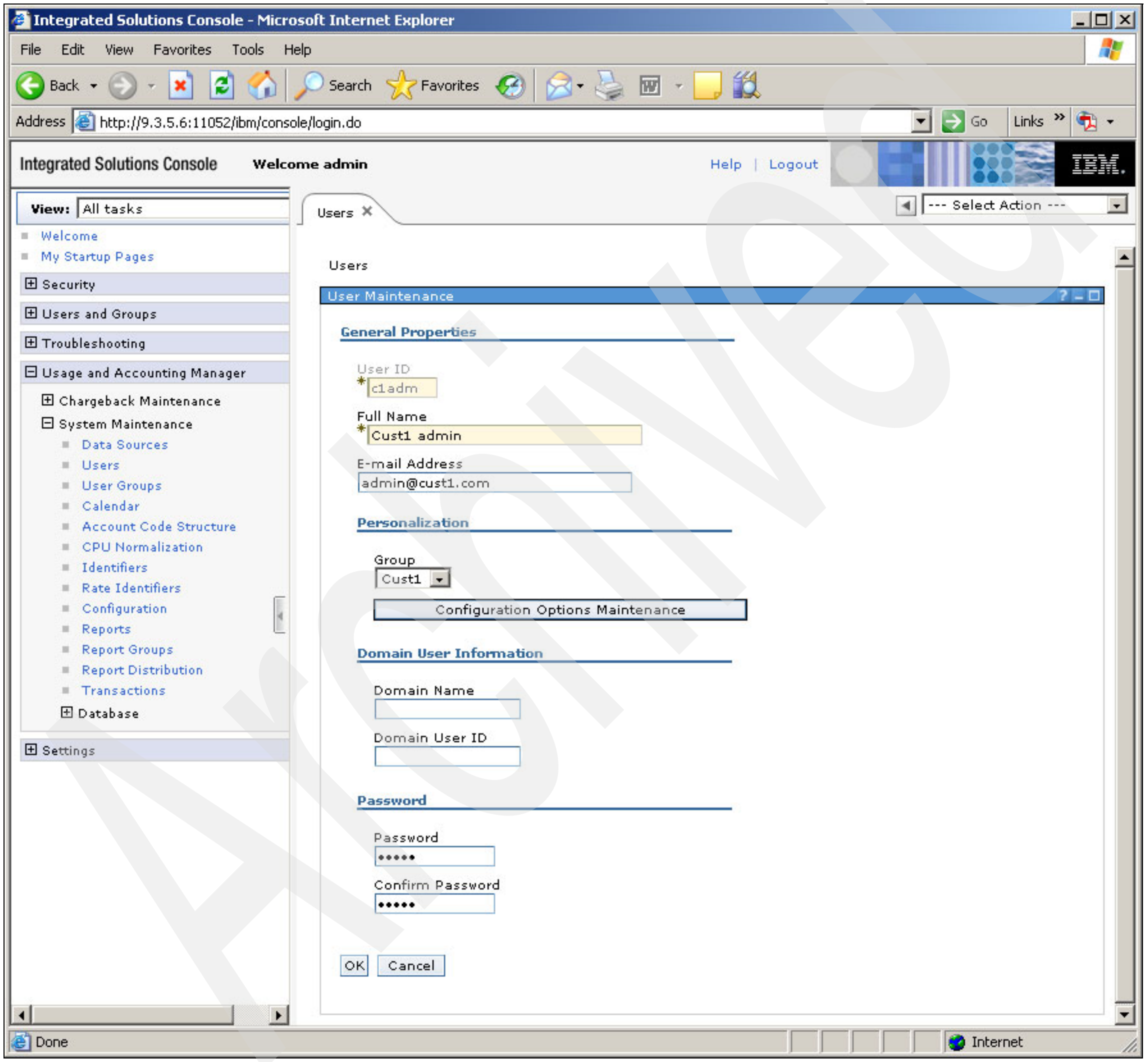

User ID: c1adm
Full Name: Cust1 admin
E-mail Address: admin@cust1.com
Group: Cust1
Domain Name:
Domain User ID:
Password: •••••
Confirm Password: •••••

Figure 5-21 Defining user

5.9.4 Tenant view

The tenant is now able to log in to the Tivoli Usage and Accounting Manager reporting server and access the reports they are authorized to access, as shown in Figure 5-22.

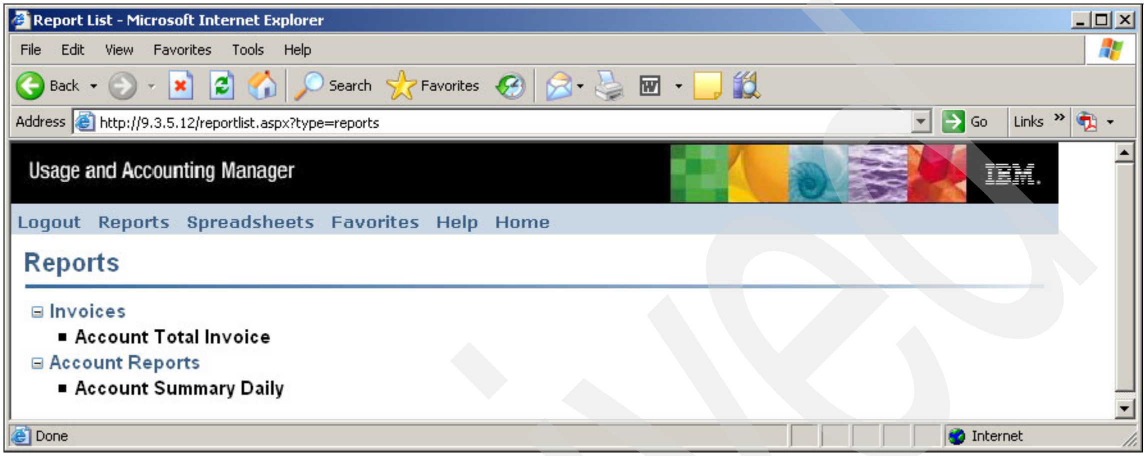
Report List - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Address http://9.3.5.12/reportlist.aspx?type=reports
Go Links
Usage and Accounting Manager
IBM.
Logout Reports Spreadsheets Favorites Help Home
Done
Internet

Reports

- Invoices
  - Account Total Invoice
- Account Reports
  - Account Summary Daily

Figure 5-22 Tenant view of Tivoli Usage and Accounting Manager reporting server

A

Software components used with TPM

This appendix indicates the software components we used to implement the provisioning with Tivoli Provisioning Manager (TPM).

Overview

To allocate and install OS with Tivoli Provisioning Manager, we used:

- Hardware Management Console (HMC)
- Tivoli Provisioning Manager (TPM)
- IBM Tivoli Open Process Automation Library (OPAL)
- Automation packages:
  - IBM Automation Package for pSeries Hardware Management Console
  - IBM Automation Package for AIX NIM installation
  - IBM Automation Package for AIX Operating System
- Additional automation packages for provisioning with TPM for AIX:
  - IBM Automation Package for AIX Logical Volume Manager
  - IBM Automation Package for AIX FIX
  - IBM Automation Package for AIX Patch Management
  - IBM Automation Package for LPAR-CPU on pSeries with AIX

Hardware Management Console

We used Hardware Management Console (HMC) Version 7 Release 3.3.0.2.

> **Attention:** Because a conflict exists between HMC Version 7 Release 3.3.0 and Tivoli Provisioning Manager Version 5.1.1.1, a safer method is to use HMC Version 7 Release 3.2.0, instead of HMC Version 7 Release 3.3.0.

IBM Tivoli Provisioning Manager

The IBM Tivoli Provisioning Manager (TPM) V5.1.1 for AIX packages are:

- IBM Tivoli Provisioning Manager for AIX, Version 5.1.1.1
- IBM Tivoli Provisioning Manager and Tivoli Intelligent Orchestrator for AIX, Version 5.1.1
- WebSphere Application Server V6.02.1.1, AIX 32 bit
- DB2 Universal Database Enterprise Server Edition for AIX, Version 8.2 FX11
- Tivoli Directory Server for AIX, Version 6.0
- NetView® Server for AIX and Solaris, Version 7.1.4

Attention: When we implemented the provisioning with TPM, we used AIX5.3TL08 instead of AIX6.1. A version conflict exists with DB2 version 8.2 and AIX6.1 during installation.

For more information about the installation of TPM, see*Tivoli Provisioning Installation Guide for AIX,* SC32-2234.

For more information about the Tivoli Provisioning Manager product, see:

- `http://publib.boulder.ibm.com/infocenter/tivihelp/v20r1/index.jsp`
- `http://`<*TPM_server_address*>`:9175/help/index.jsp`

  For <*TPM_server_address*>, enter the private network of the server itself.

IBM Tivoli Open Process Automation Library

IBM Tivoli Open Process Automation Library (OPAL) is the portal from which we can download automation libraries and best practices that use libraries. For more information about OPAL, see:

`http://www.ibm.com/software/brandcatalog/portal/opal`

IBM Automation Package for pSeries Hardware Management Console

The IBM Automation Package for pSeries Hardware Management Consolecan manage logical partitions (LPARs) on IBM POWER processor-based server under control of HMC or IVM. The package can create, delete, activate, and deactivate LPARs. It can also add or delete processors and memory from running LPARs for the dynamic allocation of resources. For information, see:

`http://www.ibm.com/software/brandcatalog/portal/opal/details?catalog.label=1TW10101G`

IBM Automation Package for AIX NIM installation

The IBM Automation Package for AIX NIM installation can manage the installation of AIX by NIM. Through using AIX NIM features, TPM can provide entire OS installation images or customized OS images with specific applications installed. For more information, see:

http://www.ibm.com/software/brandcatalog/portal/opal/details?catalog.label=1TW101005

IBM Automation Package for AIX Operating System

The IBM Automation Package for AIX Operating System can manage basic AIX commands, such as add or remove IP information, and reboot the OS. This automation package can provide a running OS that is customized for users. For rebooting, you may use other automation packages such as IBM Automation Package for pSeries Hardware Management Console. For more information, see:

http://www.ibm.com/software/brandcatalog/portal/opal/details?catalog.label=1TW10100D

Additional automation packages

Additional automation packages for provisioning with TPM for AIX include:

- IBM Automation Package for AIX Logical Volume Manager

  http://www.ibm.com/software/brandcatalog/portal/opal/details?catalog.label=1TW101025

- IBM Automation Package for AIX FIX

  http://www.ibm.com/software/brandcatalog/portal/opal/details?catalog.label=1TW10103S

- IBM Automation Package for AIX Patch Management

  http://www.ibm.com/software/brandcatalog/portal/opal/details?catalog.label=1TW10105G

- IBM Automation Package for LPAR-CPU on pSeries with AIX

  http://www.ibm.com/software/brandcatalog/portal/opal/details?catalog.label=1TW10101Q

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

For information about ordering these publications, see “How to get Redbooks” on page 279. Note that some of the documents referenced here might be available in softcopy only.

- *PowerVM Virtualization on IBM system p: Introduction and Configuration Fourth Edition,* SG24-7940
- *PowerVM Virtualization on IBM System p: Managing and Monitoring,* SG24-7590
- *IBM System p Advanced POWER Virtualization (PowerVM) Best Practices*, REDP-4194
- *IBM PowerVM Live Partition Mobility*, SG24-7460
- *Hardware Management Console V7 Handbook, SG24-7491*
- *IBM AIX Version 6.1 Differences Guide*, SG24-7559
- *Introduction to workload Partition Management in IBM AIX Version 6.1,* SG24-7431
- *Workload Partition Management in IBM AIX Version 6.1,* SG24-7656
- *AIX V6 Advanced Security Features Introduction and Configuration*, SG24-7430
- *Implementing IBM Systems Director 6.1*, SG24-7694
- *Going Green with IBM Systems Director Active Energy Manager*, REDP-4361
- *Deployment Guide Series: IBM Tivoli Usage and Accounting Manager V7.1*, SG24-7569
- *IBM Tivoli Usage and Accounting Manager V7.1 Handbook,* SG24-7404
- *Deployment Guide Series: IBM Tivoli Provisioning Manager Version 5.1*, SG24-7261
- *IBM Tivoli Application Dependency Discovery Manager Capabilities and Best Practices*, SG24-7519

- *Deployment Guide Series: IBM Tivoli Application Dependency Discovery Manager V7.1*, SG24-7616
- *Developing Workflows and Automation Packages for IBM Tivoli Intelligent Orchestrator V3.1*, SG24-6057
- *IBM Tivoli Monitoring: Implementation and Performance Optimization for Large Scale Environments*, SG24-7443
- *Deployment Guide Series: IBM Tivoli Monitoring V6.2*, SG24-7444

Other publications

These publications are also relevant as further information sources:

- Documentation available on the support and services Web site include:
  - User guides
  - System management guides
  - Application programmer guides
  - All commands reference volumes
  - Files reference
  - Technical reference volumes used by application programmers

  The support and services Web site is:

  http://www.ibm.com/systems/p/support/index.html
- Virtual I/O Server and support for Power Systems (including Advanced PowerVM feature)

  http://www14.software.ibm.com/webapp/set2/sas/f/vios/documentation/

Online resources

These Web sites are also relevant as further information sources:

- IBM AIX

  http://www.ibm.com/systems/power/software/aix/index.html
- AIX Enterprise Edition

  http://www.ibm.com/systems/power/software/aix/sysmgmt/enterprise/index.html
- Tivoli Provisioning Manager (TPM)

  http://www.ibm.com/software/tivoli/products/prov-mgr/features.htm

- Tivoli Intelligent Orchestrator (TIO)

  http://www.ibm.com/software/tivoli/products/intell-orch/
- Tivoli Application Dependency Discovery Manager (TADDM)

  http://www.ibm.com/software/tivoli/products/taddm/
- Tivoli Monitoring

  http://www.ibm.com/software/tivoli/products/monitor/
- IBM Systems Director

  http://www.ibm.com/systems/management/director/resources/index.html
- Tivoli Usage and Accounting Manager

  http://www.ibm.com/software/tivoli/products/usage-accounting/

How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

IBM

Redbooks

Multitenant Utility Computing on IBM Power Systems Running AIX

(0.5” spine)
0.475”<->0.875”
250 <-> 459 pages

Multitenant Utility Computing on IBM Power Systems Running AIX

Implement best practices for running IBM AIX in service provider environments

Leverage the latest IBM PowerVM virtualization technologies

Use management and monitoring tools provided with IBM AIX Enterprise Edition

This IBM Redbooks publication presents concepts, considerations, and high level deployment examples for technical professionals who design and operate multitenant utility computing environments hosted on IBM System p and IBM AIX Enterprise Edition. AIX Enterprise Edition brings together IBM enterprise management capabilities from Power Systems, IBM AIX, and Tivoli Software to provide a powerful set of integrated functions for infrastructure management in multitenant utility computing environments.

In this book, we focus on the following topics:

- Using System p, IBM AIX 6, and PowerVM virtualization technologies
- Provisioning in a multitenant computing environment on AIX and Power Systems
- Monitoring resources in a multitenant environment
- Accounting and chargeback in a multitenant AIX environment

AIX Enterprise Edition brings together several IBM products: AIX 6 operating system, IBM PowerVM Workload Partitions Manager, Tivoli Application Dependency Discover Manager, Tivoli Monitoring, and IBM Usage and Accounting Manager Virtualization Edition for Power Systems. In addition to these products we also provide a system allocation example using Tivoli Provisioning Manager.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-7681-00 ISBN 0738432318