# IBM

IBM AIX Enterprise Edition System Administration Guide

Learn about AIX Enterprise Edition based on AIX Version 6.1

Obtain useful planning information

Explore an offering that combines AIX with IBM Tivoli Systems Management

Shane Brandon
Erin Fitzgerald
SeongLul Son

Redbooks

ibm.com/redbooks

International Technical Support Organization

IBM AIX Enterprise Edition System Administration Guide

May 2009

First Edition (May 2009)

This edition applies to IBM AIX Enterprise Edition 6.1 (product number 5765-AEZ).

Contents

Figures

Tables

Examples

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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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

xvii

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:

AIX 5L™	POWER5™	System p®
AIX®	POWER6™	System x®
BladeCenter®	PowerHA™	Tivoli Enterprise Console®
DB2®	PowerPC®	Tivoli®
Enterprise Storage Server®	PowerVM™	TotalStorage®
Everyplace®	POWER®	WebSphere®
GPFS™	pSeries®	Workload Partitions Manager™
IBM®	Redbooks®	z/OS®
Power Systems™	Redbooks (logo) ®	
POWER4™	RS/6000®	

The following terms are trademarks of other companies:

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

Red Hat, and the Shadowman logo are trademarks or registered trademarks of Red Hat, Inc. in the U.S. and other countries.

EJB, J2EE, J2SE, Java, JDBC, JVM, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Excel, Internet Explorer, Microsoft, PowerPoint, 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 will help you install, tailor, and configure AIX® Enterprise Edition 6.1 (5765-AEZ) and its incorporated IBM Tivoli® systems management tools.

This publication provides AIX Systems Administrators a broad understanding of the new features and products that make up the AIX Enterprise Edition offering. It will help you design or create a solution to maximize efficiencies and help organizations get the most from their Power Systems™ platform.

AIX Enterprise Edition includes all the features of AIX with additional management software that is designed to complement the features of AIX V6.1. AIX Enterprise Edition is available only with AIX V6.1. AIX Enterprise Edition is designed to improve the manageability of a virtualized AIX environment. Potential benefits include:

► Improved service availability through access to relevant, real-time information about IT resources

► Enhanced operational efficiency through visualization of the relationships of IT resources and applications

► Improved operational efficiency by centralizing monitoring and automating the response to service issues

► Better utilization of IT resources through consolidated views of virtualized resource usage

► Improved service availability through Live Application Mobility

► Simplified ordering of AIX and essential service management tools at an attractive price

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, Poughkeepsie Center.

Shane Brandon is a Senior IT Specialist at IBM Australia. He has 11 years of experience in AIX and Power Systems, pSeries®, and RS/6000® hardware. He has been with IBM for five years. Shane is a Certified Advanced Technical Expert System p® as of 2006 and a Certified Systems Administrator Tivoli Storage Manager V5. His areas of expertise include Operating System

provisionng, Network Installation Management (NIM), Security, and PowerHA™. He also co-authored *Workload Partition Management in IBM AIX Version 6.1,* SG24-7656.

Erin Fitzgerald is a Senior System Administrator at IBM United States. She has 14 years of experience in AIX and Linux® systems on Power Systems and System x® hardware. She has worked at IBM for 13 years. Erin is an AIX Certified Advanced Technical Expert. Her areas of expertise include NIM, Cluster Systems Management (CSM), and security.

SeongLul Son is a Senior IT specialist working at IBM Korea. He has 13 years of experience in the IT industry and is a Certified Advanced Technical Expert for System p. His areas of expertise include networking, Power Systems virtualization, NIM, PowerHA, and General Parallel File System (GPFS™). He also co-authored the *Deploying Oracle 10g RAC on AIX 5L with GPFS,* SG24-7541, *AIX 5L Version 5.3 Differences Guide,* SG24-7463, and *AIX 5L and Windows 2000: Solutions for Interoperability,* SG24-6225.

Scott Vetter, PMP, managed the project that produced this publication.

Thanks to the following people for their contributions to this project:

Anirban Chatterjee, Jay Kruemcke, Jennifer Kuvlesky, Yoichiro Ishii
IBM, Austin

Bob Pitt and Mark Matthews
IBM, Australia

JinHyung Cho, SeHee Kim, JungJun Lim, HyunKuk Moon, and RaeSung Park
IBM Korea

Prince Ravindran
IBM India

Alex Abderrazag
IBM, UK

Henning Gammelmark
Bankdata, Denmark

David Bennin, Don Brennan, Rich Conway, Ann Lund, and Emma Jacobs
International Technical Support Organization, Poughkeepsie Center

A special thanks to Budi Darmawan of International Technical Support Organization, Austin Center. Budi's expertise and product knowledge along with his willingness to assist the team significantly contributed to and greatly aided in developing the depth and quality of this book.

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, IBM 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 IBM Redbooks publications 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

THIS PAGE INTENTIONALLY LEFT BLANK

Introduction and overview

In this chapter, we introduce AIX Enterprise Edition. We provide an overview of the individual components that make up the AIX Enterprise Edition offering and list their key features.

We describe the management functions and discuss the benefits of implementing AIX Enterprise Edition on your IBM Power Systems servers:

1.1 AIX Enterprise Edition

AIX Enterprise Edition is designed and packaged to enable administrators and organizations to simplify the managing and monitoring of their virtualized Power Systems-based environments. The AIX Enterprise Edition offering incorporates a suite of service management and monitoring tools designed to improve availability, enhance the operational efficiency of the systems within the environment, and to track, measure, and report on resource utilization and usage.

The AIX Enterprise Edition offering allows organizations to maximize all of the capabilities of virtualization, security, and performance of AIX 6.1, combined with the enterprise service management features of tools, such as Workload Partition Manager, IBM Tivoli Application Dependency Discovery Manager (TADDM), IBM Tivoli Monitoring, IBM Tivoli Usage and Accounting Manager, and DB2® Enterprise Server Edition (DB2 ESE), that allow the systems administrators to manage the environment in a more holistic way.

> **Note:** AIX Enterprise Edition is exclusive to the AIX Version 6.1 operating system. Management Edition (ME) for AIX provides a similar suite of service management tools for managing and monitoring Power Systems running the AIX Version 5.3 or AIX Version 6.1 operating system.

1.2 Comparing AIX offerings

Table 1-1 on page 3 outlines and identifies the differences among several offerings for use with AIX on the Power Systems platform. It describes what is included with each offering and the functions of the tools and products associated with them.

Table 1-1   AIX offerings and tools

Offering	Description	OS support	OS included
AIX Enterprise Edition (EE)	AIX 6 including the Workload Partition Manager, Tivoli Monitoring, Tivoli Application Dependency Discovery Manager, and Tivoli Usage and Accounting Manager Virtualization Edition	AIX 6.1	Yes
Management Edition (ME) for AIX	Tivoli Monitoring, Tivoli Application Dependency Discovery Manager, and Tivoli Usage and Accounting Manager Virtualization Edition	AIX 6.1 and AIX 5.3	No
AIX Standard Edition	The AIX operating system	AIX 6.1, AIX 5.3, and AIX 5.2	Yes
Performance Management for Power Systems	Remote performance reporting	AIX 6.1, AIX 5.3, and IBM i	No
PowerVM™ Workload Partitions Manager™ for AIX (Workload Partition Manager)	Cross-systems management for Workload Partitions and enablement for Live Application Mobility	AIX 6.1	No
IBM Systems Director	Cross-system platform management	AIX, Linux (Linux on system z, p, and x), IBM i, and Windows®	No

1.3 IBM Service Management (ISM)

Today, companies focus on providing innovative services. To deliver these services, IT and operations departments must strive to guarantee compliance, security, and continuous uptime, which all play a part in helping to ensure these business services are effectively performed to support the organization's business goals. Yet, it is common for companies with organizational silos and traditional implementations to become entrenched in managing individual areas, such as IT infrastructure technologies, single product revenues and expenses, individual processes, and organizational efficiencies, instead of managing integrated solutions and services delivered by the sum of all these components.

When these areas are isolated, there can be penalties for noncompliance and service level violations.

Enter IBM Service Management, a way to align your organization and all its related functions with your business. IBM Service Management encompasses the management processes, tactics, and best practices needed to deliver business services. IBM Service Management is about developing, deploying, and managing services, helping to reduce IT and operations costs by automating processes, and helping to more effectively manage compliance. It is about increasing flexibility and getting products, solutions, and services to market more quickly. It is about helping to respond to changes more efficiently and effectively than ever before.

IBM Service Management is designed with one thing in mind: to help you manage your business. Because IBM understands that IT and operations are very much a part of your business, we offer powerful tools to help you align the four primary components of your business:

► People
► Processes
► Information
► Technology

IBM Service Management lets you pull these critical components together with an array of tightly integrated solutions that can be viewed as three interconnected layers:

► IBM Process Management
► IBM Operational Management
► IBM Service Management Platform

1.4 Components

The AIX Enterprise Edition offering consists of a number of individual tools that combine to give an overall package designed to meet all the management and monitoring needs of the Power systems servers within the environment.

The AIX Enterprise Edition package bundle consists of:

► AIX Version 6.1

► PowerVM Workload Partitions Manager for AIX (WPAR Manager)

► IBM Tivoli Monitoring

► IBM Tivoli Application Dependency Discovery Manager (TADDM)

- IBM Tivoli Usage Accounting Manager Virtualization Edition for Power Systems
- An embedded version of DB2 Enterprise Server Edition (ESE)

  Note that this component is only offered for use with the Tivoli components

PowerVM Workload Partition Manager for AIX

Workload Partition (WPAR) Manager provides a single centralized point of control for the management of Workload Partitions in AIX 6.1. Workload Partition Manager enables Live Application Mobility, which allows for the relocation of Workload Partitions from one host to another host without disruption to the applications running in the Workload Partition.

Workload Partition Manager can be used to manage automated relocation policies for moving Workload Partitions between systems based on specified metrics, such as system load. Workload Partition manager is a browser-based interface that provides the ability for an administrator to create, clone, and remove Workload Partition definitions, or start and stop Workload Partitions from an easy to use management interface. It also contains application programming interfaces (APIs) that allow IBM Tivoli Monitoring to capture data. For further technical information about IBM PowerVM Workload Portion Manager for AIX refer to:

http://www.redbooks.ibm.com/abstracts/sg247656.html?Open

IBM Tivoli Monitoring

IBM Tivoli Monitoring for AIX is designed to monitor the health and performance of Power systems servers while having the ability to capture and analyze data. IBM Tivoli Monitoring for AIX allows a system administrator to determine trending and forecast future usage to aid capacity planning and to assist in the overall manageability of the environment.

IBM Tivoli Monitoring for AIX includes best practice solutions created by developers, which include predetermined thresholds for alerting on key metrics. The Tivoli Enterprise Portal allows an administrator a comprehensive graphical view of the state of the AIX operating system, logical partitions (LPARs), Workload Partitions, server hardware, Hardware Management Console (HMC), and Virtual I/O Server resources.

You can obtain additional product information at:

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

IBM Tivoli Application Dependency Discovery Manager

IBM Tivoli Application Dependency Discovery Manager (TADDM) uncovers application dependencies and configurations in your enterprise environment. TADDM uses the application mapping and agent-less, credential-free discovery capabilities of application dependency software to see relationships among applications, computer systems, and network devices in your environment.

TADDM discovers configurations and maps applications, changes, and address compliance measures. It can capture the configuration of each dependency, track changes to it, and provide information for analysis in order to report the history of configuration changes so that an administrator can see how configuration items have been configured and how they change over time.

TADDM compares the discovered configurations to a *reference* policy, can reveal policy violations, and determine if configurations comply with your policies. You can obtain detailed information about the IBM Tivoli Application Dependency Discovery Manager at:

http://www.redbooks.ibm.com/abstracts/sg247616.html?Open

IBM Tivoli Usage and Accounting Manager

IBM Tivoli Usage and Accounting Manager, Virtualization Edition delivers to the organization the ability to accurately track and report on usage across the Power Systems AIX environment. This product can determining resource consumption and utilization for business units, departments, or individuals in order to assist organizations in managing their assets and costs and to highlight where efficiencies can be made.

You can obtain further information at:

http://www.ibm.com/software/tivoli/products/usage-accounting

1.5 Integration of the components

AIX Enterprise Edition provides integration among the tools, such as IBM Tivoli Monitoring and IBM Tivoli Application Dependency Discovery Manager, to provide a foundation for service management across the Power Systems within your environment.

These tools are designed to deliver a set of capabilities that will assist you in getting the most from your AIX virtualized environment. AIX Enterprise Edition is designed to complement the platform management capabilities of the IBM Systems Director.

1.6 Advantages

The enterprise service management tools that make up the AIX Enterprise Edition offering have been designed make the manageability of AIX virtualized environments more efficient and to allow organizations to maximize the performance and utilization of their Power Systems environments.

The AIX Enterprise Edition provides visualization of system resources and application utilization and their behavior either in a historical view or in real time. It provides for a single central interface for monitoring and automation of all aspects of the environment and the ability to manipulate systems within the environment as well as to monitor the usage and accounting from all AIX hosts in the environment.

Live Application Mobility and the centralization function of the Workload Partition Manager provide for a single point of control in the environment for managing Workload Partitions (WPARs). Workload Partition Manager for AIX provides for improved system utilization and application performance by enabling more efficient load balancing throughout the operating system and enterprise environment. It improves application availability by giving the systems administrator the flexibility to seamlessly relocate workloads from one host to another host without disruption to the application, thus, increasing availability by reducing the need for planned outages due to system maintenance.

1.7 Hardware prerequisites

AIX Enterprise Edition has no prerequisites beyond the prerequisites of AIX Version 6.1 and can, therefore, be utilized on any hardware platform that supports the AIX 6 Operating System, including:

► IBM pSeries POWER4™ technology-based servers
► IBM System p POWER5™ technology-based servers
► IBM Power Systems POWER6™ technology-based servers
► IBM BladeCenter® JS21 with PowerPC® 970 processors
► IBM BladeCenter JS22 with POWER6 processors

> **Note:** For the latest supported servers, see:
>
> http://www-03.ibm.com/systems/power/software/aix/sysmgmt/enterprise/requirements.html

1.8 Lab overview

The ITSO infrastructure setup for this project, which is seen in Figure 1-1, consisted of:

► IBM System p 550 (9113-550)

► Two Power Systems 570s (9117-MMA)

► A Hardware Management Console

► IBM TotalStorage® DS4800 (1815-88A)

The System p 550 was configured as a Network Installation Management (NIM) server and provided a centralized management role for our environment. Each of the Power Systems 570 servers were configured with a Virtual I/O Server, four partitions (one for each of the Tivoli products), and a generic AIX partition that hosted several Workload Partitions for both application and system.

Figure 1-1   ITSO project infrastructure

Planning and sizing

In this chapter, we discuss the planning and sizing considerations for the installation and configuration of AIX 6.1 Enterprise Edition. We look at the requirements for planning the installation of each of the Tivoli tools included with Enterprise Edition, as well as Workload Partition Manager. We also provide a brief overview of logical partitioning (LPAR) and Workload Partition sizing:

9

2.1 Planning and sizing virtualization

Virtualization enables the sharing or the aggregation of physical resources and the creation of logical resources that are used by AIX V6.1 Enterprise Edition. AIX V6.1 Enterprise Edition is not 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 into a larger one, which is similar to joining several small disks in one large virtual disk. Also. virtualization can provide for uniform access to resources.

2.1.1 Virtualization advantages

Several advantages of virtualization are:

► Improved hardware utilization

 Server hardware resources can be more 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 infastructure

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

► Rapid provisioning

 Virtualization can enable rapid infastructure 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 can be deployed 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 the overall power usage and cooling requirements.

2.2 Planning and sizing logical partitions

Each logical partition (LPAR) must have a set of resources available. Shared processor LPARs, also called *micro-partitions*, enable you to partition a processor into multiple virtual partitions and provide the flexibility to change the allocation of system resources dynamically.

With the launch of POWER6 processor-based servers, IBM added key features, such as the ability to migrate a running LPAR between systems (PowerVM Live Partition Mobility). The minimum LPAR resources that are needed are:

► At least 0.1 unit of processor for micro-partitions

► At least 256 MB of main memory

► At least one disk to store the operating system (for AIX, the rootvg)

► At least one disk adapter or integrated adapter to access the disk

► At least one LAN adapter per partition to connect to the Hardware Management Console (HMC)

► A partition must have an installation method, such as Network Installation Management (NIM), and a means of running diagnostics, such as network diagnostics

2.3 Planning and sizing Workload Partition Manager

AIX workload partitions (WPARs) provide a software-based virtualization solution for creating and managing multiple AIX operating system environments within a single AIX-based logical partition. The WPAR provides isolation of software services, applications, and administration within a single instance of the AIX V6.1 operating system. WPARs can be used on any server platform that is capable of running AIX V6.1 Enterprise Edition, including POWER4 technology-based servers and beyond.

When building a WPAR from the command line, you can configure and start it within a few minutes:

► Network File System (NFS) is a prerequisite to the WPAR mobility functionality.

► At least 256 MB of main memory is required.

► At least one disk to store the operating system (for AIX, the rootvg) is required.

► At least one disk adapter or integrated adapter to access the disk is required.

► At least one LAN adapter per partition to connect to the HMC is required.

2.3.1 Additional resources for virtualization

There are several additional IBM Redbooks publications about virtualization, as well as LPAR mobility. Several additional resources on these topics are:

► *IBM PowerVM Live Partition Mobility*, SG24-7460

► *Multitenant Utility Computing on IBM Power Systems Running AIX*, SG24-7681

► *PowerVM Virtualization on IBM System p: Introduction and Configuration Fourth Edition*, SG24-7940

► *Workload Partition Management in IBM AIX Version 6.1*, SG24-7656

2.4 Planning and sizing DB2

The two major considerations when planning the installation of DB2 Enterprise Server Edition are the space for the package images and then the space required for the database. In addition to the space requirements, you must create users for the database administration and instance. We discuss each of these items in this section.

Note that the DB2 Enterprise Server Edition as part of AIX Enterprise Edition is an embedded part of the Tivoli components and is not a separately licensed product.

2.4.1 DB2 installation requirements for DB2 server

You must meet the following operating system, hardware, and communications prerequisites to install a DB2 server product (Table 2-1 on page 13).

Table 2-1   Installation requirements for DB2 Server

AIX Version	Operating system	Hardware
6.1	64-bit AIX kernel xlC.rte 9.0.0.0 xlC.aix61.rte 9.0.0.0	All processors that are capable of running the supported AIX 6.1 operating system: Power 4 or above

2.4.2 DB2 file systems

Planning for installing DB2 includes considering the creation of additional file systems and, in certain cases, additional Volume Groups, depending on the environment on which you install DB2. The minimum DB2 installation requires:

► /opt/IBM/db2

Approximately 600 MB contains the DB2 source code installed on your system. This file system is static and does not grow after the initial installation.

► /home/db2inst1

This name is the default name for the database instance home. It varies in size depending on database configuration and grows over time as data is added to the database.

> **Note:** The name /home/db2inst1 can be any file system name. In many installations, the administrator chooses a name for the file system that is the same name as the application. Therefore, in our example, /home/db2inst1 is `/tuam_data/ituamdb`. We recommend that you create these file systems on a separate AIX Volume Group and not on the rootvg.

You need to consider other DB2 file systems, such as the DB transactional logs and the database backups. We recommend that these file systems are broken out into their own file system for better database performance when the database grows large. For more information about configuring and setting up DB2, refer to *DB2 Deployment Guide*, SG24-7653.

2.4.3 DB2 users and groups

Setting up user IDs and groups for DB2 is similar to setting up user IDs and groups for other applications in AIX. Table 2-2 on page 14 provides a list of the DB2 default user IDs and the groups that are created as part of the base installation of DB2.

Table 2-2   Default DB2 user IDs and groups

User ID	Group
dasusr1	dasadm1
db2fenc1	db2fgrp1
db2inst1	db2grp

These IDs can be used to create your database instance and install your Tivoli
application, or you can create a custom DB2 user ID and group to help identify
the instance with which the db2 user ID is associated. Example 2-1 shows how to
create a DB2 user ID and group specifically for Tivoli Usage and Accounting
Manager.

Example 2-1   Custom DB2 user ID

```
TUAM10:]mkgroup -A uamdbgrp
TUAM10:]mkuser -a "pgrp=uamdbgrp" ituamdb
TUAM10:]mkgroup -A uamdbadm
TUAM10:]mkuser -a "pgrp=uamdbadm" dbadmusr
TUAM10:]mkgroup -A uamfgrp
TUAM10:]mkuser -a "pgrp=uamfgrp" uamfusr
```

2.5 Planning and sizing IBM Tivoli Monitoring V6.2

In the following sections, we describe the hardware requirements for the IBM
Tivoli Monitoring V6.2 infrastructure for distributed AIX systems. We describe the
following components:

- ► Hub monitoring server
- ► Remote monitoring server
- ► Portal server
- ► Portal client
- ► Tivoli Data Warehouse
- ► Warehouse Proxy agent
- ► Summarization and Pruning agent

2.5.1 IBM Tivoli Monitoring processor considerations

For the best performance, we recommend that you use multiprocessor systems in the following scenarios:

► Running Tivoli Enterprise Portal client on a computer that is also running one of the server components.

► Monitoring the environment of 1000 or more monitored agents with multiple server components on the same computer. For example:

– Portal server and hub monitoring server

– Monitoring server (hub or remote) and Warehouse Proxy agent

– Warehouse Proxy agent and Summarization and Pruning agent

► Running a small environment with all the server components (monitoring server, portal server, Warehouse Proxy agent, and Summarization and Pruning agent) on a single computer.

► For the Tivoli Data Warehouse database server. Except for a small installation, use a multiprocessor system for the Tivoli Data Warehouse database. If you install the Warehouse Proxy agent on the warehouse database server, consider using a two-way or four-way processor.

► If you install the Summarization and Pruning agent on the warehouse database server (with or without the Warehouse Proxy agent), consider using a four-way processor. For large environments where more CPU resources might be needed, we can run the Summarization and Pruning agent on a computer separate from the warehouse database server. In this case, ensure that a high-speed network connection exists (100 Mbps or faster) between the Summarization and Pruning agent and the database server.

2.5.2 IBM Tivoli Monitoring memory and disk sizing requirements

Table 2-3 on page 16 shows the estimated memory and disk storage requirements for IBM Tivoli Monitoring components on distributed systems.

Table 2-3   Memory and disk requirements

Component	Memory requirements[a]		Disk storage requirements[b]
	Small environment[c]	**Large environment[d]**	
Hub Monitoring Server	70 MB	100 MB	650 MB[e]
Remote Monitoring Server	100 MB	300 MB	250 MB[e]
Portal Server	100 MB[f]	300 MB[f]	800 MB
Portal Client (Browser or Desktop)	150 MB	300 MB	150 MB
Tivoli Data Warehouse	2-4 GB depending on database configuration paramenters	2-8 GB depending on database configuration paramenters	Follow the *IBM Tivoli Monitoring Version 6.2.0 Installation and Setup Guide*, GC32-9407 to estimate the database size.[g]
Warehouse Proxy agent	50 MB	100 MB	150 MB
Summarization and Pruning agent	100 MB	300 MB	150 MB

a. The memory and disk sizings shown in this table are the amounts required for the individual component beyond the needs of the operating system and any concurrently running applications.
b. The disk storage estimates apply to any size monitoring environment and are considered high estimates. The size of log files affects the amount of storage that is required.
c. A small environment is considered to be a monitoring environment with 500 to 1000 agents, with 100 to 200 monitored agents per remote monitoring server.
d. A large environment is considered to be a monitoring environment with 3000 or more monitored agents, with 500 to 1000 monitored agents per remote monitoring server.
e. The storage requirements for the hub and remote monitoring servers do not include storage for the agent depot, which can require an additional 1 GB or more.
f. The memory requirement for the portal server does not include database processes for the portal server database, which require up to 400 MB of additional memory, depending on configuration settings.
g. One of the factors to consider when planning the size of the database that you need is the amount and type of information that you will collect for agent history data collection. Refer to the "Planning Considerations for the Tivoli Data Warehouse" chapter in the *IBM Tivoli Monitoring Version 6.2.0 Installation and Setup Guide*, GC32-9407, to estimate database size.

2.5.3 IBM Tivoli Monitoring database estimation

The primary considerations when sizing the database are the amount and the type of data that will be collected and the type of historical data that will be retained. Each agent user's guide provides capacity planning information to help

a system administrator calculate the amount of disk space that is required of the data for each attribute group. This information can then be used to complete the following calculations that will determine how large the database needs to be:

► Determine the number of detailed records per day for each attribute group

► Determine the storage footprint for each attribute group

► Determine the amount of detailed data for the attribute group

► Calculate the amount of aggregate data for the attribute groups

Note: IBM Tivoli Monitoring is capable of recording large amounts of data if the recording interval is short, and if the scope of the recording is large. Validate your estimates with testing when your installation is complete.

The system administrator might prefer to use the planning speadsheet available in the IBM Tivoli Monitoring information center that is shown in Figure 2-1. Alternatively, you might prefer to use the Warehouse Load Projection tool that is available at the IBM Tivoli Open Process Automation Library at the following URL:

http://www.ibm.com/software/tivoli/opal

This tool does all the calculations for you and includes data for the monitoring agents.

| Attribute group | Number of agents | Data from the User's Guide | | | Collection interval | Detailed records per day* | Attribute agent disk space (KB)* | Days of detailed data kept | Warehouse space - detailed (MB)* | Aggregate records | Warehouse space - aggregated (MB)* | Total warehouse space for attribute group (MB) |
		Record size - detailed	Record size- aggregated	Interval instances								
Total warehouse data size												
(sum of all attribute group total warehouse space)												
Total database size												
(Total warehouse data size * 1.5)												

*:
　Use the equations in Estimating the required size of your database for these values.

Figure 2-1   Tivoli data warehouse database size estimation worksheet

2.5.4 IBM Tivoli Monitoring server protocols and values

You must consider which of the four possible communications protocols to use for your environment. The differences are distinguished in Table 2-4.

Table 2-4   Tivoli Monitoring server protocols

Protocol	Value	Definition
IP.UDP	IP.UDP port number. Default is 1918.	The port number for the monitoring server. The default is 1918.
SNA	Net name	The SNA network identifier for your location.
	Logical unit (LU) name	The LU name for the monitoring server. This LU name corresponds to the Local LU Alias in your SNA communications software.
	Log mode	The name of the LU6.2 LOGMODE. The default value is CANCTDCS.
IP.PIPE	IP.PIPE port number. Default is 1918.	The port number for the monitoring server. The default is 1918.
IP.SPIPE	IP.SPIPE port number. Default is 3660.	The port number for the monitoring server. The default is 3660.

2.5.5 IBM Tivoli Monitoring additional considerations

There are no additional software requirements beyond the elements packaged as part to the AIX Enterprise Edition offering, such as Java™ and the Xlc runtime environment that is also required to support the DB2 instance required by IBM Tivoli Monitoring.

A video card supporting 64 000 colors and 1024 x 768 resolution is required for the portal client.

2.6 Planning and sizing TADDM

In this section, we explain several terms related to TADDM and show an architecture of TADDM. In addition, we provide the planning and sizing considerations for hardware.

2.6.1 TADDM terminology and architecture

In TADDM, a *domain* is a logical subset of a company's infrastructure. It can be based on any convenient boundary criteria, such as geographical or departmental. Domains can delineate organizational (for example, IT operations or line of business (LOB) operations), functional (for example, finance IT domain or human resources IT domain), or geographical (for example, United States IT domain or Brazil IT domain) boundaries. There is typically one TADDM Server deployed per domain to discover domain applications and their supporting infrastructure, configurations, and dependencies.

The TADDM Server is also referred to as the *Domain Manager*. You might want to use one or multiple TADDM Servers (Domain Managers), depending on how many configuration items (CIs) exist in your environment. If you have a large environment with thousands of hardware, software, and other configuration items, one Domain Manager might not be enough, and you can expand to multiple Domain Managers.

You can have a view of all of the aggregated data at an *Enterprise Configuration Management Database (eCMDB) Domain Manager*. In an eCMDB environment, the participating Domain Managers perform the discoveries and supply data to the enterprise Domain Manager. The enterprise Domain Manager aggregates information from the Domain Servers. This behavior is configurable depending on the needs of an organization. Figure 2-2 on page 20 shows the relationship between Enterprise Domain Manager and Domain Manager.

Figure 2-2   Multiple Domain Servers

Discovery *sensors* reside on the TADDM server and collect configuration attributes and dependencies. TADDM offers a wide variety of discovery sensors to enable immediate discovery of virtually all components that are found in the typical data center, across the application software, host, and network tiers.

Entities discovered by TADDM

TADDM discovers configurations and interdependencies across the following environment entities:

► Application components, such as Web servers, application servers, and databases

► System components, such as hosts, operating systems, load balancers, and database servers

► Network components, such as routers, switches, and firewalls

► Services, such as Domain Name System (DNS) and Lightweight Directory Access Protocol (LDAP) services

Table 2-5 on page 21 shows entities discovered by TADDM in detail.

Table 2-5   Entities discovered by TADDM

Entity	Description
Network tier	TADDM discovers the following devices at the network tier, along with the MIB2 (RFC 1213) parameter values for each device: ▶ Routers ▶ Switches ▶ Load balancers ▶ Firewalls ▶ Generic IP devices
System tier	TADDM discovers the following devices at the system tier: ▶ Server hosts and disks ▶ Host IP interfaces ▶ Database servers ▶ Load balancers or clusters
Application tier	TADDM discovers the following components at the application tier, along with version, configuration files and properties, host information, and vendor-specific extensions for each component (except the generic processes): ▶ Custom servers, based on custom templates that you design ▶ Java 2 Platform, Enterprise Edition (J2EE™) application servers and configurations ▶ J2EE and Java 2 Platform, Standard Edition (J2SE™) components and modules ▶ Web server components ▶ Web modules, configuration files, and installation directories ▶ Generic Java virtual machine (JVM™) processes ▶ Databases
Infrastructure service components	The system infrastructure services that support the application environment are discovered along with the dependency structure between these service components and the application components. The infrastructure service components are: ▶ DNS and Network File System (NFS) services ▶ LDAP

Entity	Description
Relationship structure	In addition to the discovery of components, the physical and logical connectivity at the network, system, and application tiers is discovered at the following level of support in each of the tiers: ► Layer 3 IP connectivity ► Layer 2 connectivity ► Application component runtime dependencies ► Infrastructure service dependencies

Additional Required Software for TADDM

In order to install TADDM and to discover components at a later time, the following additional software is required:

► Database server is required for TADDM. You can use any existing databases or newly install the DB2 bundled in the IBM Management Edition (ME) for AIX.

► On a TADDM server:

  – Firefox if you plan to use Launchpad for installing ME for AIX
  – unzip
  – sudo
  – ssh
  – lsof
  – X filesets

► On target servers:

  – ssh
  – lsof

> **Note:** You can install the ssh filesets from the Expansion Pack, while the unzip, sudo, and lsof files are in AIX Toolbox for Linux.

2.6.2 TADDM hardware sizing

This section describes the general sizing rules of TADDM and database.

TADDM server

The TADDM server must have the following attributes:

► 100 GB of available disk space
► 2 - 4 CPUs with a minimum process speed of 2 GHz
► 4 - 8 GB of memory
► 4 - 8 GB of swap space

> **Note:** The size of a server that is shown in this section handle up to hundreds or thousands of target servers, depending on how many CIs each target server has. However, if you do not have many target servers, a small amount of CPU and memory might be suitable in your environment. For example, we have the TADDM server and database installed on the same server with 1 CPU (4.2 Ghz) and 6 GB of memory in our test environment.

It is difficult to provide a standard number of CIs in a server. The actual number of CIs in a server varies depending on the complexity of the infrastructure; for example, a complex database server with a large number of instances, databases, and table spaces has a larger number of CIs per server, and the number of CIs affects the overall performance. For further information about sizing TADDM, refer to the following publication: *Deployment Guide Series: IBM Tivoli Application Dependency Discovery Manager V7.1*, SG24-7616.

TADDM database sizing considerations

To calculate a database size, use this general rule: 2 MB per target host. For example, the following numbers of hosts work out to the following database sizes:

5000 hosts   10 GB

10 000 hosts   20 GB

An actual large-scale implementation had the following numbers. Having been running weekly discoveries of about 2500 hosts and network devices for 14 months uses about 7 GB of space. The reason that the planning information cannot be better qualified is that the determining factor, besides the number of hosts and devices, is the number of components that are discovered in the environment and the rate of change. Both of these factors can vary tremendously from one environment to another environment. In addition, the size of an individual component varies as well, for example, how many EJBs are deployed in a WebSphere® Application Server, or how many tables exist in a database? Using Custom Server Templates and Computer System Templates to collect and track files can also greatly influence the capacity requirements of the CMDB.

2.7 Planning and sizing IBM Tivoli Usage and Accounting Manager

IT departments are faced with unparalleled pressures to operate as businesses and align IT costs with company priorities, which is difficult to accomplish, especially when those resources are shared in your infastructure. With Tivoli

Usage and Accounting Manager, Virtualization Edition (VE), you can help solve this problem by collecting, tracking, and reporting this usage. Understanding the overall usage of the system gives you a better picture to more effectively utilize CPU and memory resources in a shared environment, which can help manage costs.

> **Note:** IBM Tivoli Usage and Accounting Manager, Virtualization Edition (VE) as part of AIX Enterprise Edition or Management Edition for AIX is fully supported.

2.7.1 IBM Tivoli Usage and Accounting Manager architecture

Tivoli Usage and Accounting Manager is comprised of several major components: the administration application, the reporting component, and the processing component. All of these components connect to a data source to store and retrieve information. There are several possible configurations that you can use. You need to evaluate your requirements and resource availability to determine the best configuration for your environment.

IBM Tivoli Usage and Accounting Manager components

This section covers the components for Tivoli Usage and Accounting Manager:

► Administration server

The administration server is based on the Integrated Solutions Console (ISC). ISC is an administration tool that runs inside an embedded WebSphere Application Server. The administration server component for Tivoli Usage and Accounting Manager, Virtualization Edition, is installed on the AIX V6.1 TUAM server box that will run DB2.

► Processing server

The data processing server is typically run on the same machine as the administration server. However, there might be a requirement, such as load balancing, that precludes running the data processing server on the same machine as the administration server. The Tivoli Usage and Accounting Manager processing engine applies business rules to the data and stores information in a database. The component will also run on AIX V6.1 Enterprise Edition.

► Reporting server

For Tivoli Usage and Accounting Manager, Virtualization Edition, the reporting is supported via Business Intelligence and Reporting Tools (BIRT). BIRT is a tool that runs in the Eclipse environment and can be installed on Windows, Linux, or AIX. BIRT needs to be downloaded from the

http://www.eclipse.org Web site and installed separately from IBM Tivoli Usage and Accounting Manager.

► Database server

The data server stores the information that is produced by the processing engine. It does not have to run on the same system or platform as the processing engine. DB2 V9.1 is supported for Tivoli Usage and Accounting Manager and is included with the installation CDs. You can select to install DB2 at the time that you select to install IBM Tivoli Usage and Accounting Manager.

Figure 2-3 shows the Tivoli Usage and Accounting Manager Virtualization Edition layout. In the next sections, we describe the components of IBM Tivoli Usage and Accounting Manager and the database to help you plan for the installation in your environment.
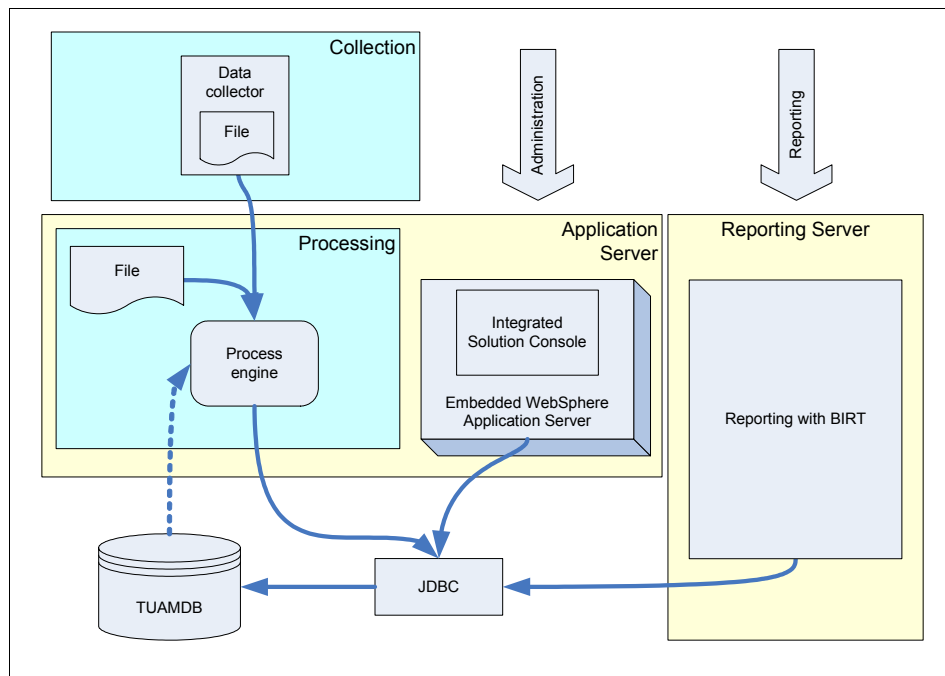
Figure 2-3   IBM Tivoli Usage and Accounting Manager architecture

2.7.2 IBM Tivoli Usage and Accounting Manager requirements

Table 2-6 on page 26 provides a list of the minimum IBM Tivoli Usage and Accounting Manager software requirements.

Table 2-6   Minimum software requirements for IBM Tivoli Usage and Accounting Manager

Operating system	AIX 5.3 or AIX 6.1
Browser	Firefox 2
Hard disk drive space	5 GB minimum, 40 GB recommended (available hard disk drive space) Note: This size varies by environment.
Processor speed	3 GHz minimum
Memory	2 GB minimum
Core software	► If you use DB2 and the server is on another server than the database, you must install the DB2 client connection software on the TUAM server. ► The Java Database Connectivity (JDBC™) drivers for DB2. The drivers are installed with DB2 and are db2jcc.jar and db2jcc_license_cu.jar. ► Red Hat® Package Manager (RPM) ► Embedded WebSphere Application Server (installed during IBM Tivoli Usage and Accounting Manager installation) ► Integrated Solutions Console (ISC) installed during the IBM Tivoli Usage and Accounting Manager installation ► BIRT Report Engine and Report Designer available for download from http://www.eclipse.org for running data usage reports.

Table 2-7 provides a list of the core software installation files.

Table 2-7   Core requisite file

Required file	Contents
EmbeddedExpress_aix_ppc32.tar	Embedded WAS server (eWAS)
ISCA71_4_EWASv61.tar	Integrated Solutions Console (ISC)

2.7.3 Estimating IBM Tivoli Usage and Accounting Manager database size

In this section, we discuss estimating the Tivoli Usage and Accounting Manager database growth. Because it is a data and collection processing tool, it collects

and loads data into the database and keeps it for a certain period of time. Estimating its growth is critical to ensuring that the space is properly located and the resulting performance impact can be addressed (such as the time to back up the data, query response time, response needs, and so on.)

We start by checking our database size in the AIX file system directly after creation. The data size is roughly 350 MB, including the database catalog and the database log files.

Data elements

The primary growth in the database comes from the following usage and accounting data:

Resource utilization   The collection of the resource usage metrics from the AcctCSR file. Collection is provided by an identifier for each resource (rate code). This collection is optional. You do not need to collect the resource usage.

Billing summary   This data provides a summary for each resource (rate code) by account code. This data links to the identifier table for to get the identifier key for each of the entries here. The data is one-to-one mapping for the BillDetail.txt file.

Billing detail   This data provides individual entries from the AcctCSR file. It gives the individual occurrences of source usage by resource name (rate code). This data links to the identifier table to get the identifier key for each of the entries here. The data is one-to-one mapping from the Ident.txt file.

Identifier table   This data lists the identifiers that are used by each billing detail entry. The data is one-to-one mapping from the Ident.txt file.

Important tips to keeping the database size manageable are:

► Run the DBpurge program using the Job runner to remove old data. Because Tivoli Usage and Accounting Manager data is an accounting financial tool, you might want to archive the data first.

► Sort by the Account_Code field before running the billing, or use a Common Source Resource (CSR) or Common Source Resource Plus (CSR+) file for billing input.

► Only collect the identifiers and resources in which you are interested. Modify the sample collection jobs, change the mapping, and remove any unwanted identifiers and resources fields.

IBM Tivoli Usage and Accounting Manager database growth factors

The growth factors for the database are:

Number of days   The retention of your data before you run the purge step to remove it.

Number of shifts   The number of shifts in a day that needs unique rate codes.

Collection source   Each collection is processed with various jobs. Each collection will generate a unique set of data.

Number of identifiers Each identifier is put in a separate row in the CIMSDETAILIDENT table.

Identifier mix   This number is the number of unique identifiers in each collection. You must be able to estimate this number based on your understanding of the collection process.

Table 2-8 provides a list of how the database rows map to the database table.

Table 2-8   Table estimation and mapping

Name	Row size	Affecting source
CIMSRESOURCE UTILIZATION	300	Source, Account_Code, Identifier mix, RateCode, Shift, and number of days
CIMSSUMMARY	300	Source, Account_Code, RateCode, Shift, and number of days
CIMSDATAIL	350	Source, Account_Code, Identifier mix, Rate per id, Shift, and number of days
CIMSDETAILIDENT	75	*<Ident mix>* x *<Ident count>*

3

Installation and configuration

In this chapter, we provide information about installing and configuring each product that is included in AIX Enterprise Editions (EE):

3.1 Launchpad

Launchpad is a program that helps you to install all the Tivoli products included in EE for AIX from a centralized GUI. EE for AIX consists of Launchpad, DB2, and three Tivoli products. The list of media shipped in the EE for AIX package is provided in Table 3-1.

Table 3-1   The list of media that is included in EE for AIX

Product	Media format	Approximate size
Launchpad V2.2	DVD	476 MB
DB2 ESE V9.1 Fix Pack 4	DVD	326 MB
IBM Tivoli Monitoring V6.2 Fix Pack 1 Tools	DVD	2.66 GB
IBM Tivoli Monitoring System p agents	CD	312 MB
TADDM V 7.1	DVD	4.17 GB
IUAM VE7.1	CD	555 MB

Installing Firefox

In order to use the Launchpad, you need to install a Web browser on AIX. We installed Firefox browser on our system. Example 3-1 on page 31 shows the steps for installing Firefox on AIX.

We downloaded the Firefox code from the following Web site:

http://www-03.ibm.com/systems/power/software/aix/browsers/index.html

The prerequisite files are also included in the downloaded filesets. Install those files first, and then install Firefox.

> **Note:** When installing the Red Hat Package Manager (RPM) prerequisites for Firefox, if you get an error message about missing X11 libraries (libX11.a(shr.o) and so forth) that are already installed on your system, you need to run the following command:
>
> ```
> # /usr/sbin/updtvpkg
> ```
>
> On AIX, the RPM package database creates a virtual package named `AIX-rpm`, which represents a snapshot of all of the `installp` files installed on the system. However, this virtual package can become out-of-date if you install additional `installp` software on the system after you have installed the `rpm.rte` fileset. The **/usr/sbin/updtvpkg** command will update the virtual package `AIX-rpm` to match the current state of the system.

Example 3-1   Installing Firefox on AIX

```
## Download and unzip Firefox.base.tar.gz

taddm10:/mnt/Firefox ] ls
Firefox                           gtk2-2.8.3-9.aix5.1.ppc.rpm
Firefox.base.tar.gz               libjpeg-6b-6.aix5.1.ppc.rpm
Java5_64.sdk.tar.gz               libpng-1.2.8-5.aix5.1.ppc.rpm
atk-1.10.3-2.aix5.1.ppc.rpm       libtiff-3.6.1-4.aix5.1.ppc.rpm
cairo-1.0.2-6.aix5.1.ppc.rpm      pango-1.10.0-2.aix5.1.ppc.rpm
expat-1.95.7-4.aix5.1.ppc.rpm     xcursor-1.0.2-3.aix5.1.ppc.rpm
fontconfig-2.2.2-5.aix5.1.ppc.rpm xft-2.1.6-5.aix5.1.ppc.rpm
freetype2-2.1.7-5.aix5.1.ppc.rpm  xrender-0.8.4-7.aix5.1.ppc.rpm
gettext-0.10.40-6.aix5.1.ppc.rpm  zlib-1.2.3-3.aix5.1.ppc.rpm
glib2-2.8.1-3.aix5.1.ppc.rpm

taddm10:/mnt/Firefox ] rpm -ivh *.rpm
error: failed dependencies:
        libX11.a(shr4.o) is needed by cairo-1.0.2-6
        libX11.a(shr4.o) is needed by gtk2-2.8.3-9
        libXext.a(shr.o) is needed by gtk2-2.8.3-9
        libX11.a(shr4.o) is needed by pango-1.10.0-2
        libX11.a(shr4.o) is needed by xcursor-1.0.2-3
        libX11.a(shr4.o) is needed by xft-2.1.6-5
        libX11.a(shr4.o) is needed by xrender-0.8.4-7

taddm10:/mnt/Firefox ] ls -l /usr/sbin/updtvpkg
-rwxr-x---   1 root     system         3880 Mar  2 17:11 /usr/sbin/updtvpkg

taddm10:/mnt/Firefox ] /usr/sbin/updtvpkg
Please wait...

taddm10:/mnt/Firefox ] rpm -ivh *.rpm
```

```
taddm10:/mnt/Firefox/Firefox ] ls
.toc                      README.install
Firefox.base              download-prerequisites.sh

taddm10:/mnt/Firefox ] smitty installp

                        Install Software

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

[TOP]                                              [Entry Fields]
* INPUT device / directory for software           .
* SOFTWARE to install                             [Firefox.base      > +
  PREVIEW only? (install operation will NOT occur)  no                +
  COMMIT software updates?                          yes               +
  SAVE replaced files?                              no                +
  AUTOMATICALLY install requisite software?         yes               +
  EXTEND file systems if space needed?              yes               +
  OVERWRITE same or newer versions?                 no                +
  VERIFY install and check file sizes?              no                +
  Include corresponding LANGUAGE filesets?          yes               +
  DETAILED output?                                  no                +
  Process multiple volumes?                         yes               +
  ACCEPT new license agreements?                    yes               +
[MORE...8]

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

If you have successfully completed the installation of the Web browser, put the Launchpad media into the media drive and run `launchpad.sh`. Remember that you need to have a graphical user environment to invoke the Launchpad with the Web browser.

Running the `launchpad.sh` file prompts you to accept the license agreement as shown in Figure 3-1. Read the license agreement for ME for AIX (the same panel will be displayed for the EE installation), and click **Accept**.
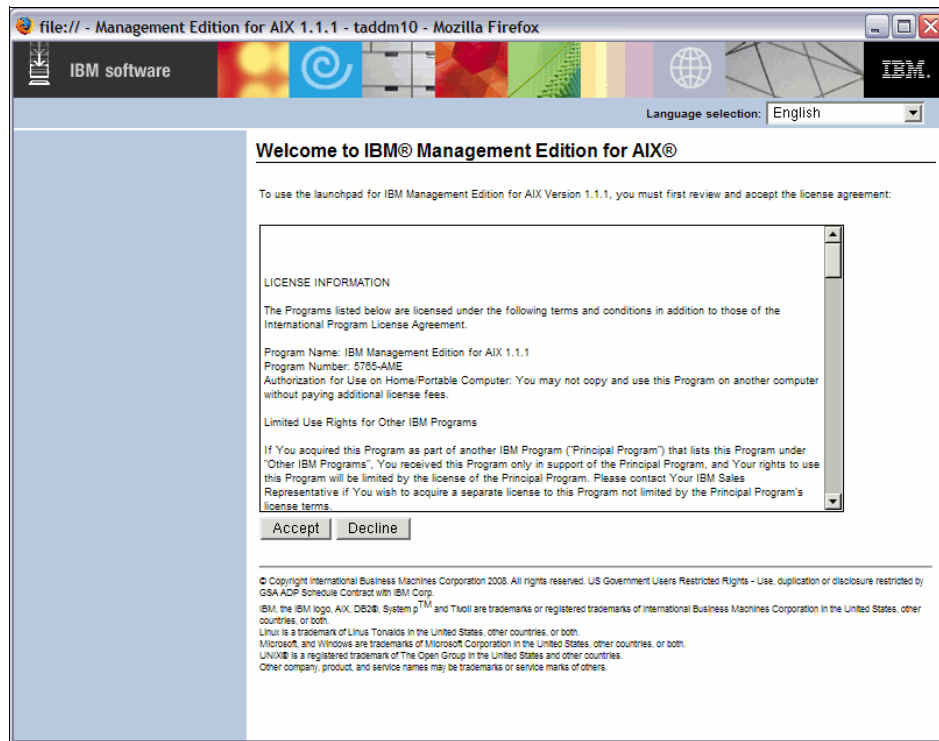

Figure 3-1   License agreement for ME for AIX

Figure 3-2 shows the main page of the Launchpad where you can install all the products in ME for AIX. Browse each menu in the left pane for future installation.
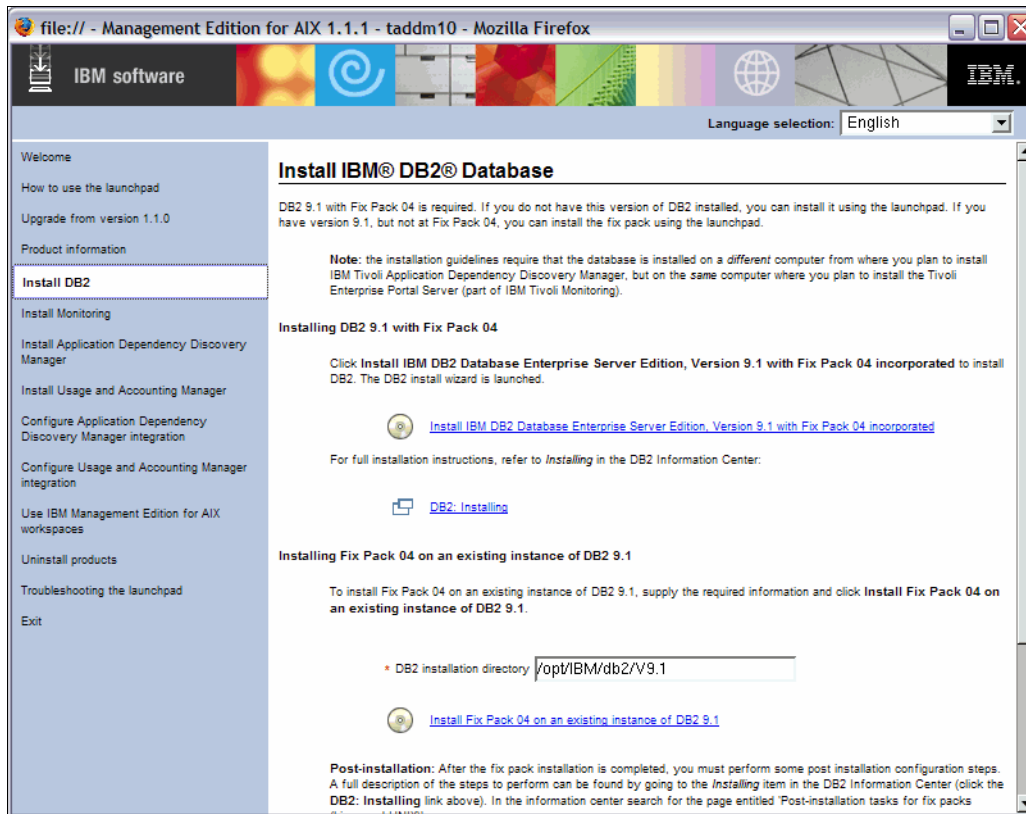

Figure 3-2   The main page of the Launchpad

3.2 Installing and configuring DB2

For our implementation, we chose to use DB2 for the database. In this section, we describe how to install DB2 using the Launchpad. Before installing DB2, we created separate file systems for the DB2 source code and the database instance home on a non-root volume group as shown in Example 3-2 on page 35:

/opt/IBM/db2   Approximately 600 MB contains the DB2 source code installed on your system. This file system is static and does not grow after the initial installation.

/home/db2inst1   This name is the default name for the database instance home. It varies in size depending on the database configuration, and it grows over time as data is added to the database.

Example 3-2 shows the installation and configuration of DB2 for the TADDM server that is used in this publication. You can select other naming conventions that are appropriate for your environment.

Example 3-2   File systems created for DB2

```
taddm10:/ ] lspv
hdisk0          00c1f170aa7f638b                        rootvg          active
hdisk1          00c1f170b4c335e3                        appvg           active
taddm10:/ ] lsvg -l appvg
appvg:
LV NAME               TYPE       LPs     PPs     PVs  LV STATE      MOUNT POINT
loglv00               jfs2log    1       1       1    closed/syncd  N/A
fslv00                jfs2       128     128     1    closed/syncd
/home/db2inst1
fslv02                jfs2       128     128     1    closed/syncd  /opt/IBM/db2
taddm10:/ ] mount /home/db2inst1
taddm10:/ ] mount /opt/IBM/db2
```

Install DB2 Enterprise Server

To simplify the installation process, we decided to use the Launchpad that is part of ME for AIX to install DB2. Launchpad is software that helps you install all the products in ME for AIX through the Web GUI interface.

We obtained the Launchpad and DB2 media and copied the install images into our Network File System (NFS) server. From the NFS server, we then exported the file systems and then mounted those file systems containing the install images on our AIX host, `taddm10`. We used the following directories for the NFS mount points:

/mnt/launchpad   For the Launchpad media

/mnt/db2   For the DB2 media

To install DB2 Enterprise Server:

1. Log on to the `taddm10` server as root and change the directory to where the Launchpad is located, which is `/mnt/launchpad`.

2. Run the `launchpad.sh` command to display the Management Edition (ME) for AIX Launchpad panel, which is shown in Figure 3-3 on page 36.

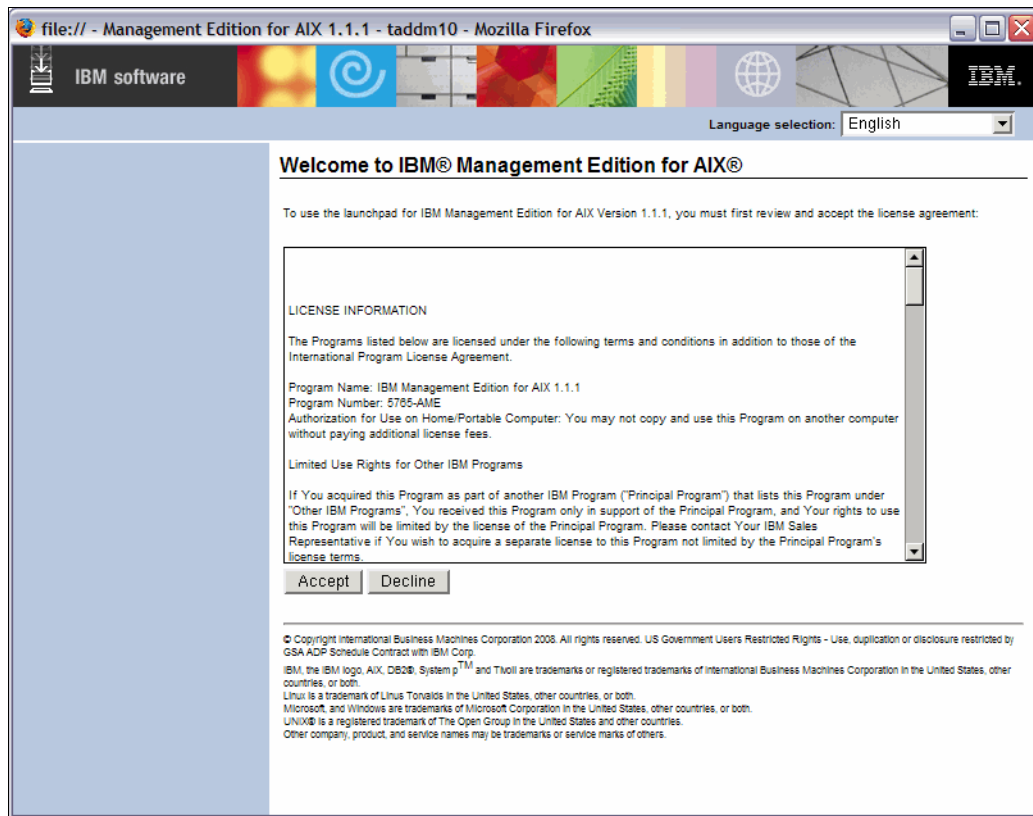3. Read the license agreement for ME for AIX and click **Accept**.

Figure 3-3   License agreement for ME for AIX

The next window (Figure 3-4 on page 37) shows the main page of Launchpad where you can install all the products in ME for AIX.

4. Select **Install DB2** from the left pane.

There are two options that you can choose: you can either install DB2 with the fix pack or you can just apply the fix pack on an existing instance.

5. Select **Install IBM DB2 Database Enterprise Server Edition, Version 9.1 with Fix Pack 04 incorporated** as shown in Figure 3-4.
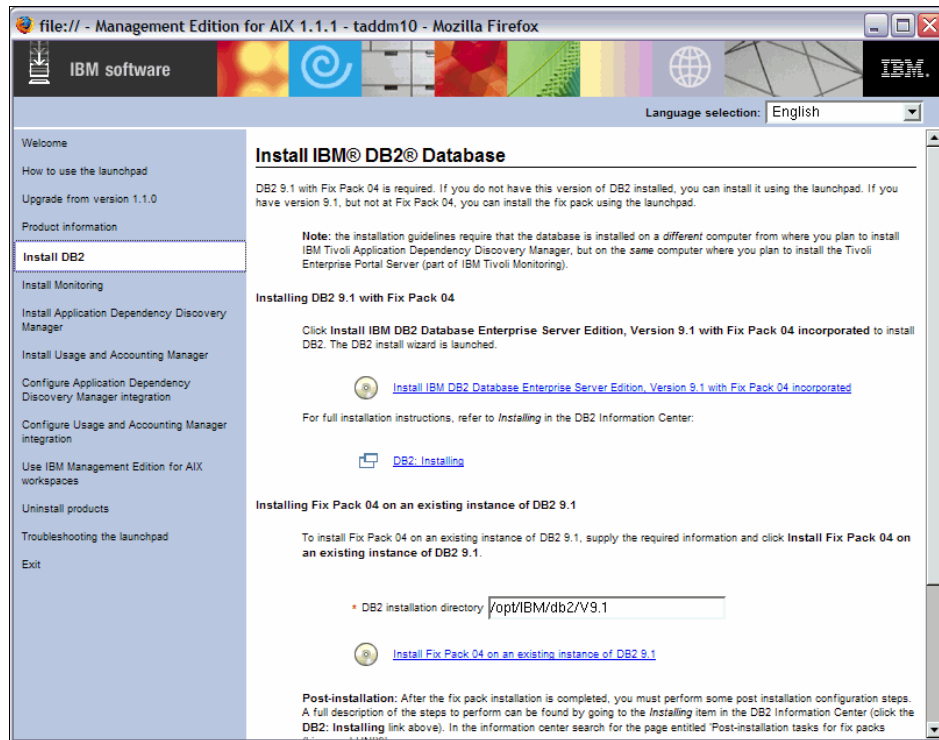

Figure 3-4   Installing DB2 through Launchpad

6. In the next window, provide the installation setup directory path where the DB2 installation image resides as shown in Figure 3-5. Click **OK**.


Figure 3-5   Selecting DB2 setup directory

7. Click **Install New** to install DB2 as shown in Figure 3-6. We need to have at least 120 MB of free space for the `/tmp` file system.

Figure 3-6 Installing a product panel

8. The next window shows the Welcome to the DB2 Setup wizard panel, which is shown in Figure 3-7.

   The DB2 Setup wizard leads you through the installation process. Notice the list of 14 installation steps on the left side of the wizard panel. As the setup wizard leads you through the installation of DB2, this list reminds you where you are in the installation process.
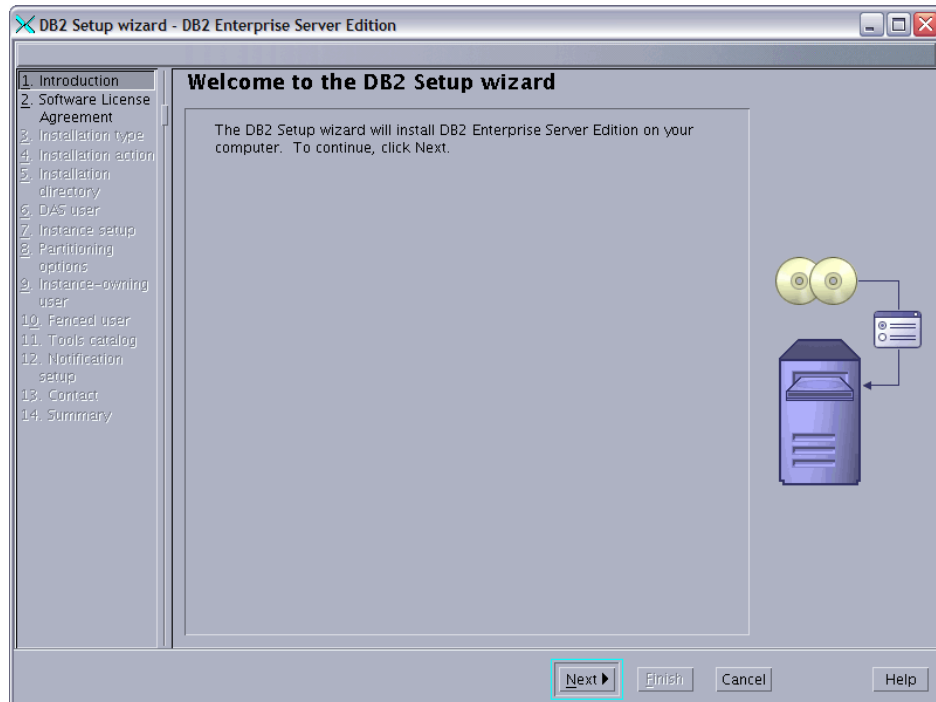
9. Click **Next.**

Figure 3-7   Welcome to the DB2 Setup wizard

The next window is the Software License Agreement panel. Read through the licensing agreement.

10. If you agree with the terms, click **Accept** and then click **Next** as shown in Figure 3-8.
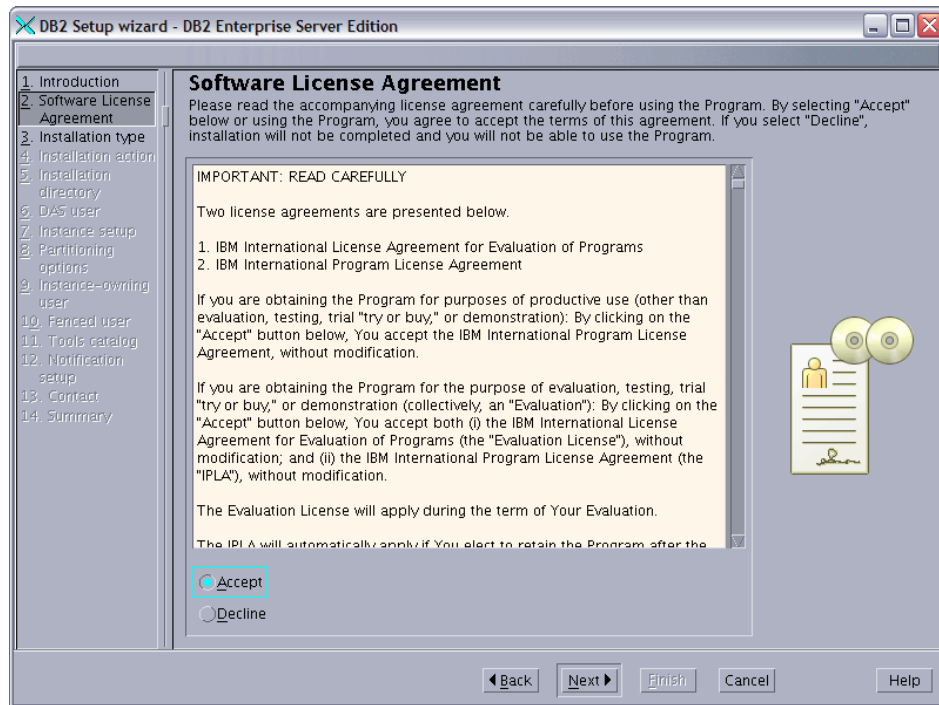

Figure 3-8   Software License Agreement

11. The next window shows the Select the installation type panel, which is shown in Figure 3-9. Select **Typical: 490 - 590 MB** and then click **Next**.
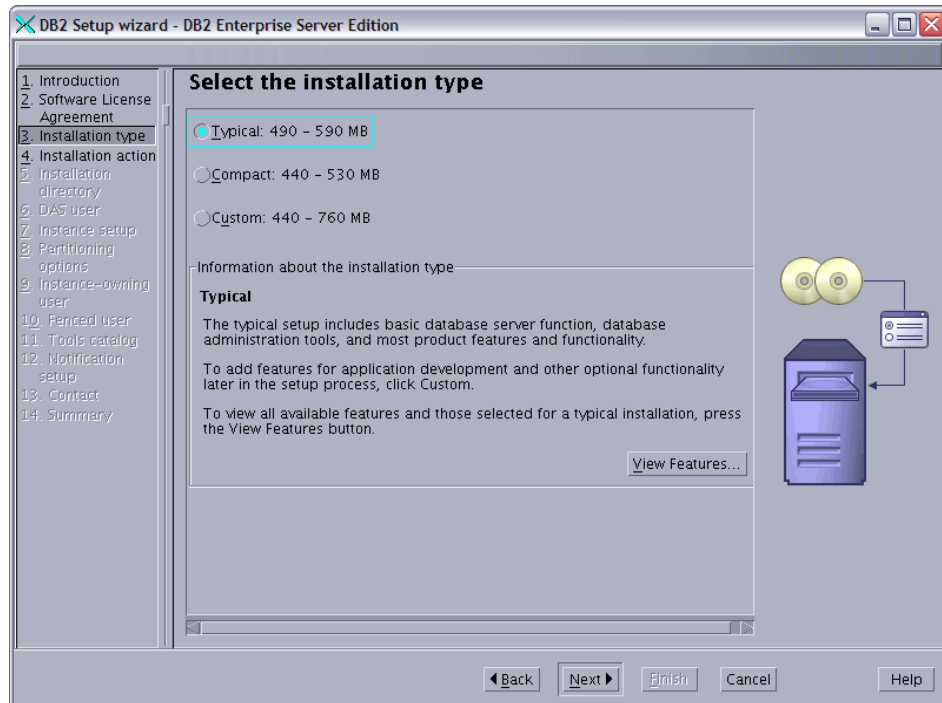

Figure 3-9   Selecting the installation type

In the next window, you have an option to choose whether to create a response file with the same configuration for future installation, which is shown in Figure 3-10.

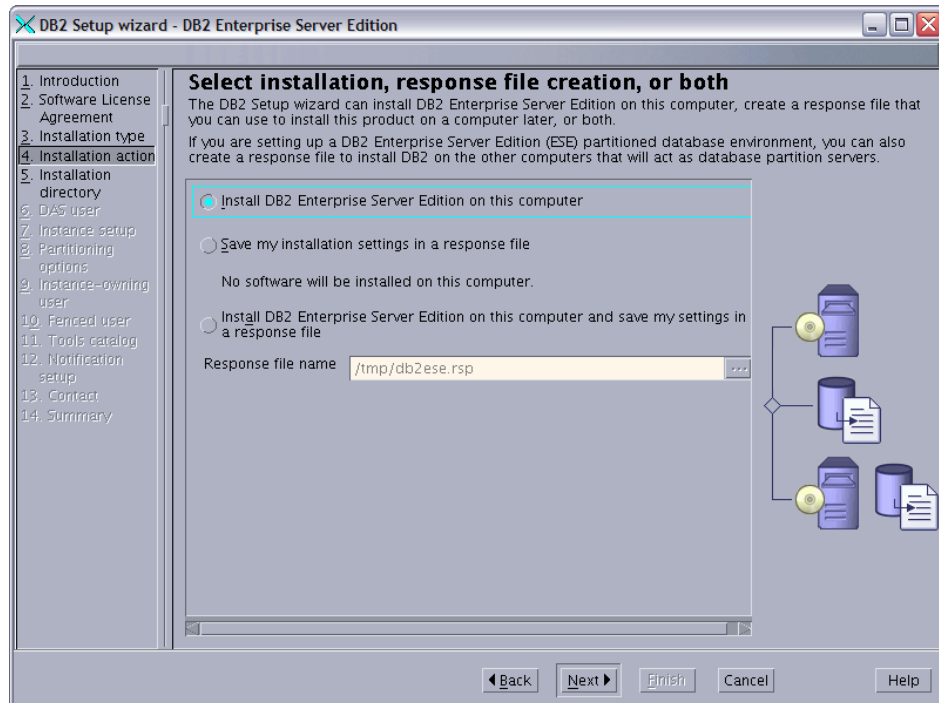12. Select **Install DB2 Enterprise Server Edition on this computer** and click **Next**.

Figure 3-10   Selecting install, response file creation, or both panel

In the next window, you select the installation directory. Provide a directory where you want to put the DB2 files.

13. For a basic install, select the default and click **Next**, which is shown in Figure 3-11.
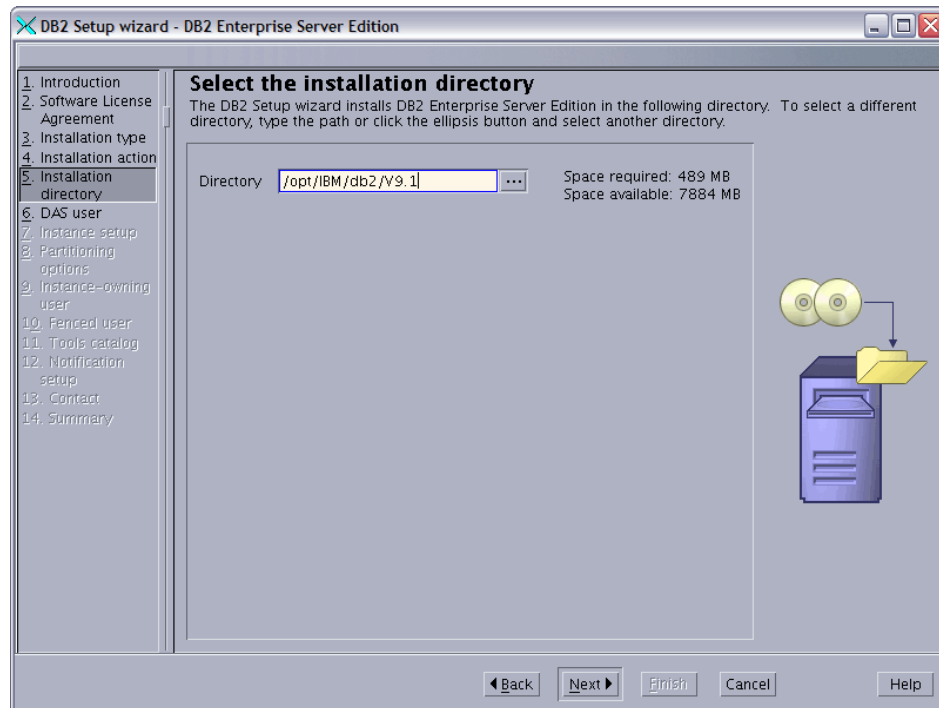

Figure 3-11   Select the installation directory

14. On the next window, enter the information for the DB2 Administration Server (DAS) user information.

   We accepted the default values and entered a password for the user, which is shown in Figure 3-12.

15. Click **Next**.

Figure 3-12   Set user information for the DB2 Administration Server

The next window lets you create a DB2 instance. In DB2, an *instance* is an independent environment where database objects can be created and applications can be run against them.

16.We will create the DB2 instance later; therefore, select **Do not create a DB2 instance** as shown in Figure 3-13.

17.Click **Next**.

> **Important:** If you operate in an environment that requires a consistent user ID/group ID (UID/GID) for applications, database users, and groups, you need to manually assign and manage these UIDs/GIDs to ensure consistency between hosts in your environment. Installing via the GUI, by default, will allocate the next available ID, which might not be the most appropriate in all environments. In addition, in case you set up an expiration policy for a password, set up database users so that they cannot be expired.
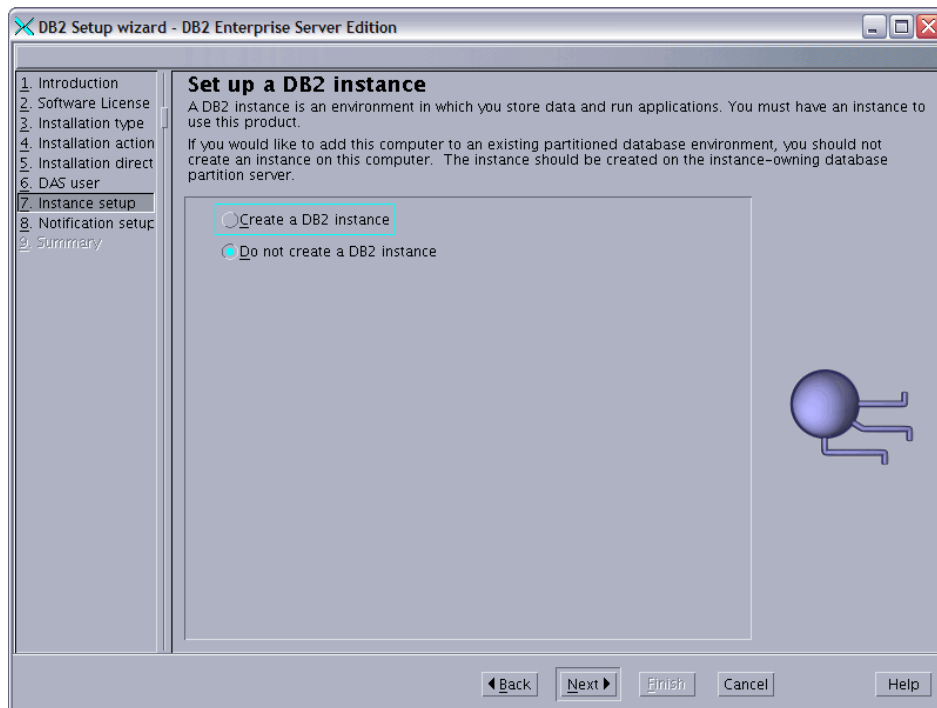

Figure 3-13   Set up a DB2 instance

In the next window, you can set up notifications.

18. Select **Do not set up your DB2 server to send notifications at this time** as shown in Figure 3-14.
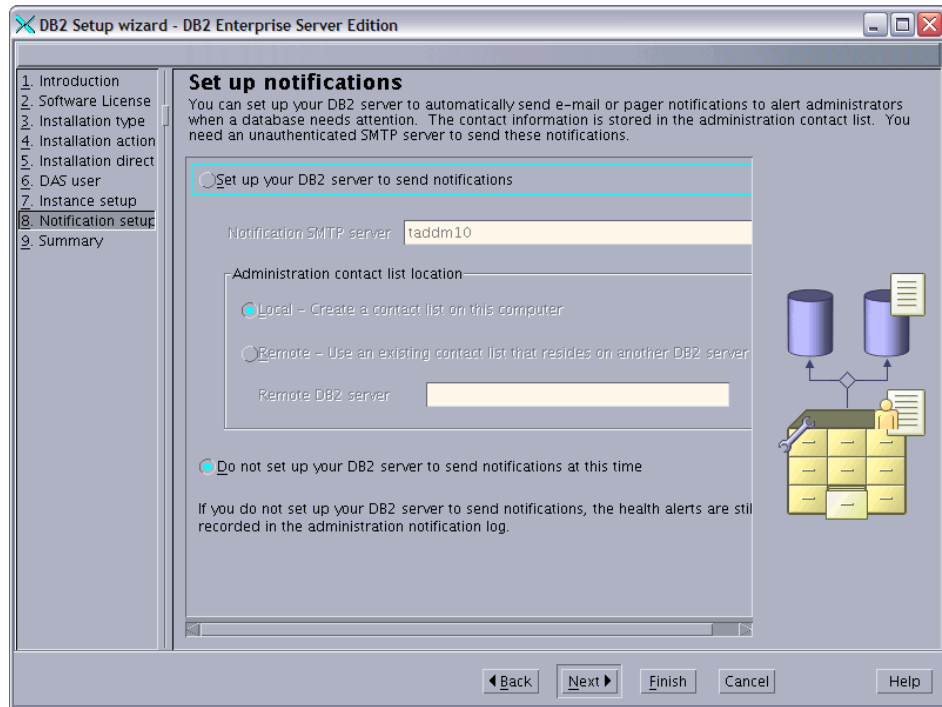
19. Click **Next**.

Figure 3-14   Set up notifications

In the next window, you can review the current settings.

20.If you are satisfied with the settings, then click **Next** as shown in Figure 3-15.
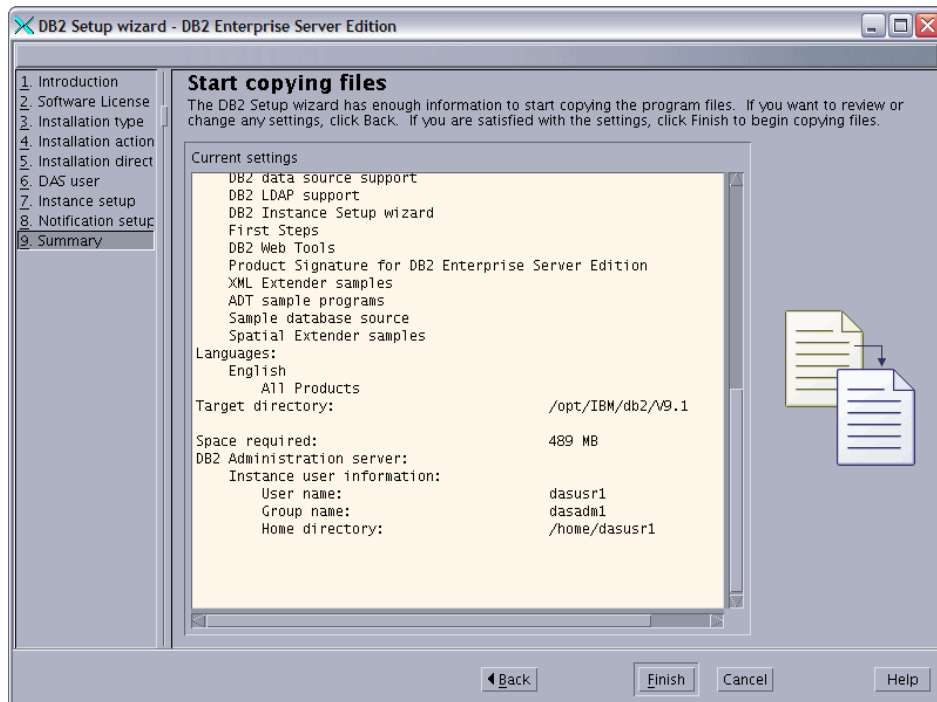

Figure 3-15   Start copying files

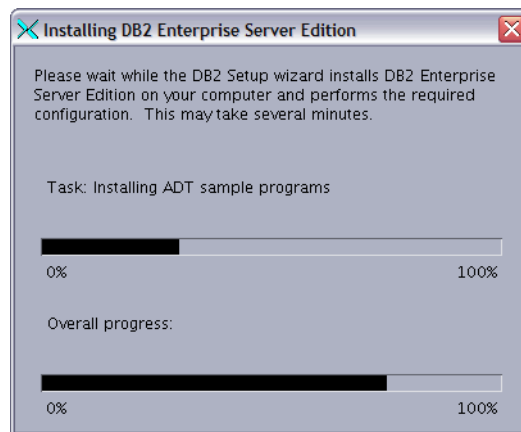In the next window, a progress panel appears while the files are copied as shown in Figure 3-16.

Figure 3-16   installation progress panel

The next window (shown in Figure 3-17) is displayed when all of the files have been copied and DB2 Enterprise Server is successfully installed.
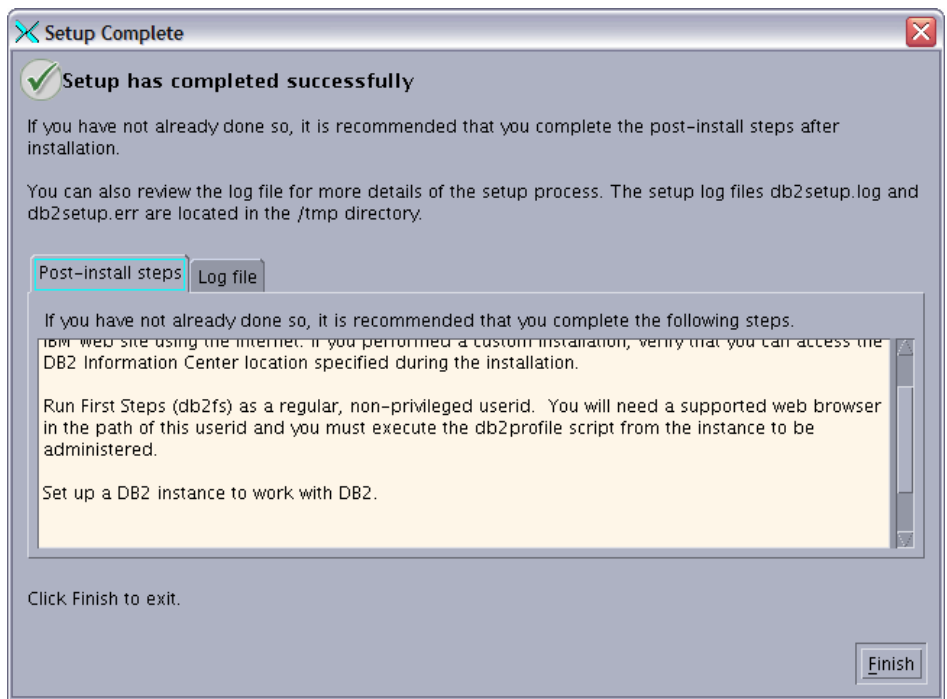
Figure 3-17   Setup has completed successfully

Create DB2 database users

Now that we have installed DB2 Enterprise Server, we need to create the database instance. We created two user accounts before we create the database instance. These accounts, which are also known as the DB2 user and the DB2 archuser, are the primary and secondary users for the database instance, as shown in Table 3-2.

Table 3-2   DB2 database users

TADDM environment (host name)	Primary/ DB2 user	Secondary/ DB2 Archuser	DB2 instance name
taddm10	db2inst1	archuser	cmdb

Example 3-3 on page 49 shows the AIX commands that create the required users.

Example 3-3   Creating users for db2 instance

```
taddm10:/mnt/TADDM71 ] mkgroup db2grp1
taddm10:/mnt/TADDM71 ] mkuser pgrp=db2grp1 db2inst1
taddm10:/mnt/TADDM71 ] mkuser pgrp=db2grp1 archuser
taddm10:/mnt/TADDM71 ] passwd db2inst1
Changing password for "db2inst1"
db2inst1's New password:
Re-enter db2inst1's new password:
taddm10:/mnt/TADDM71 ] passwd archuser
Changing password for "archuser"
archuser's New password:
Re-enter archuser's new password:
taddm10:/mnt/TADDM71 ] login db2inst1
db2inst1's Password:
[compat]: You are required to change your password. Please choose a new one.

db2inst1's New password:
Re-enter db2inst1's new password:

$ login archuser
archuser's Password:
[compat]: You are required to change your password. Please choose a new one.

archuser's New password:
Re-enter archuser's new password:$
```

Create the DB2 instances

Next, you need to create the DB2 instance. We changed the directory and
created the instance as shown in Example 3-4.

Example 3-4   Creating DB2 instance

```
taddm10:/ ] chown db2inst1.db2grp1 /home/db2inst1
taddm10:/ ] ls -ld /home/db2inst1
drwxr-xr-x    3 db2inst1 db2grp1          256 Mar  3 14:38 /home/db2inst1

taddm10:/opt/IBM/db2/V9.1/instance ]db2icrt -a SERVER -p 50010 -s ese -u
db2inst1 db2inst1    <
DBI1070I Program db2icrt completed successfully.
```

3.3 Installing and configuring IBM Tivoli Monitoring V6.2

In order to step through the installation of IBM Tivoli Monitoring V6.2, we build a lab environment, as shown in Figure 3-18. Here, we can see the topology using one of product features of the Tivoli Enterprise Portal: the self-monitoring topology.
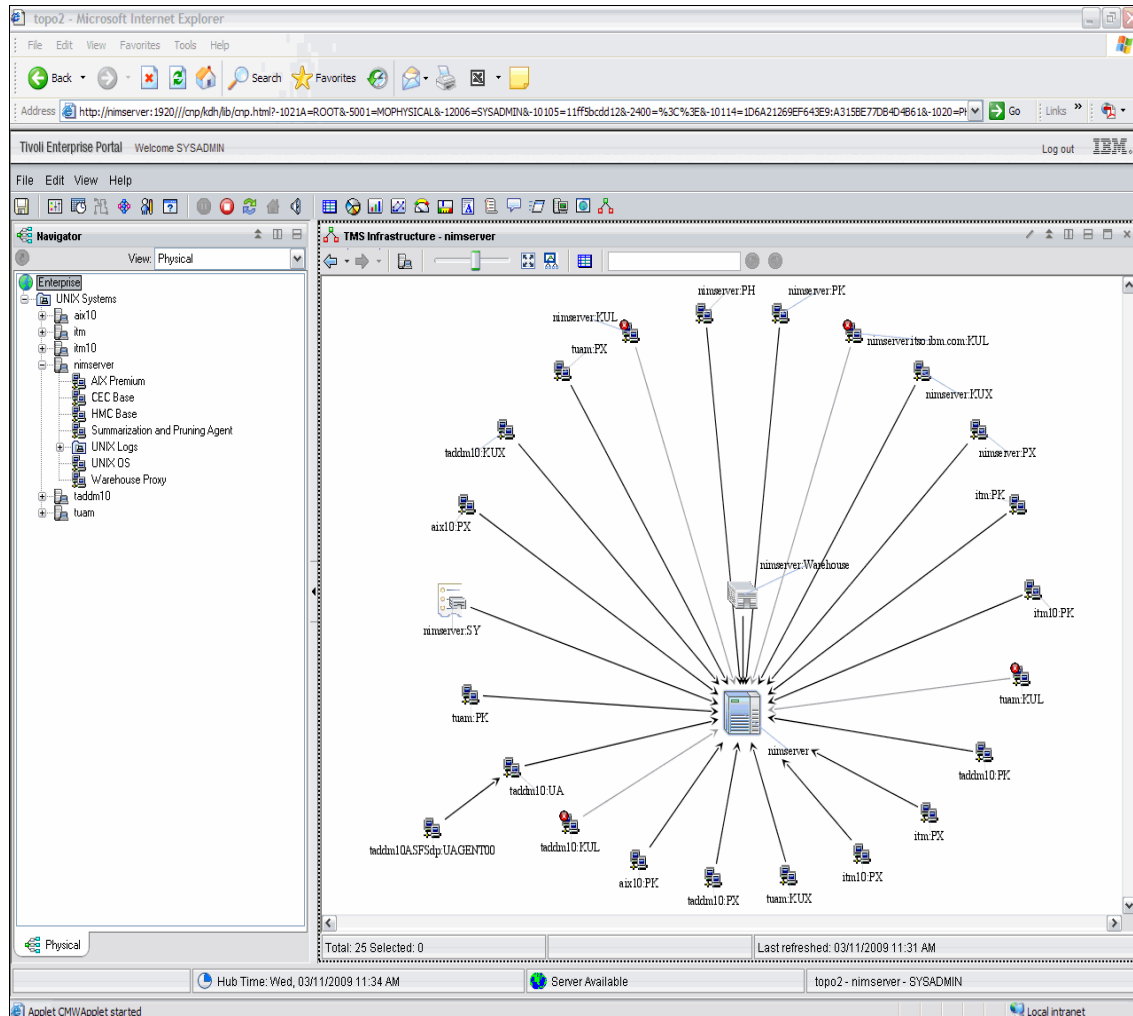

Figure 3-18   Lab environment self-monitoring topology

3.3.1 Pre-installation preparation

Before you begin, take note of the following information concerning the installation procedures in this section.

The installation procedures provide information for installing a single component (such as the monitoring server) on one computer. If you want to install multiple components (such as the monitoring server and the portal server) on the same computer and you want to install them simultaneously using the command line, the actual steps might vary. Refer to the *IBM Tivoli Monitoring Version 6.2.0 Installation and Setup Guide*, GC32-9407, for further information.

You need to consider several basic principles and guidelines when determining the procedure and setup that will best suit your environment.

For example, you can install a monitoring server on one host, a portal server on another host, and access the portal using a Windows desktop client.

If your site uses the IP.PIPE protocol for communications, be aware of the following considerations:

► There can be at most 16 IP.PIPE processes per host.

► IP.PIPE uses one, and only one, physical port per process. Port numbers are allocated using a well-known port allocation algorithm. The first process for a host is assigned port 1918, which is the default.

► KDC_PORTS is not supported for IP.PIPE.

► If you need to have more than 16 processes per host, use IP.UDP (User Datagram Protocol) for connections between IBM Tivoli Monitoring V6.2 components.

► Determine whether your environment requires a remote Tivoli Enterprise Monitoring Server hub or data warehouse.

We also need to plan for settings that will vary from one environment to another environment, such as:

► Host names

► UIDs and GIDs for the administration of IBM Tivoli Monitoring

► File systems and sizing

► Port numbers (unless the default port numbers will be used)

► Whether the monitoring server being installed or being connected to is configured as a hub or as a remote monitoring server

► Whether you want to configure a hot standby

3.3.2 Installing IBM Tivoli Monitoring from the Launchpad

The AIX Enterprise Edition Launchpad runs on IBM AIX 6.1 at a minimum technology level 1. The Launchpad runs in a Web browser window. Supported browsers are Mozilla Firefox for AIX, Versions 1.5.0.10, 1.5.0.11, and 1.5.0.12. Refer to the following Web site for a download of the last named version and information about its prerequisites:

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

Unless you are installing IBM Tivoli Monitoring by using the Launchpad on the console, you will also require the presence of a running X Windows emulator, because certain parts of the install spawn a shell for certain responses that determine the configuration. In our environment, we used Hummingbird Exceed 2008.

The prerequisites of the products that the Launchpad installs differ from the prerequisites for the browser. For details, consult the Product Information page in the Launchpad, which provides the documentation of the supported versions of those products and links to further reference material located in the IBM Tivoli Information Center.

To install IBM Tivoli Monitoring, from the LaunchPad initial window that is seen in Figure 3-19 on page 53, select **Install Monitoring** from the sidebar.
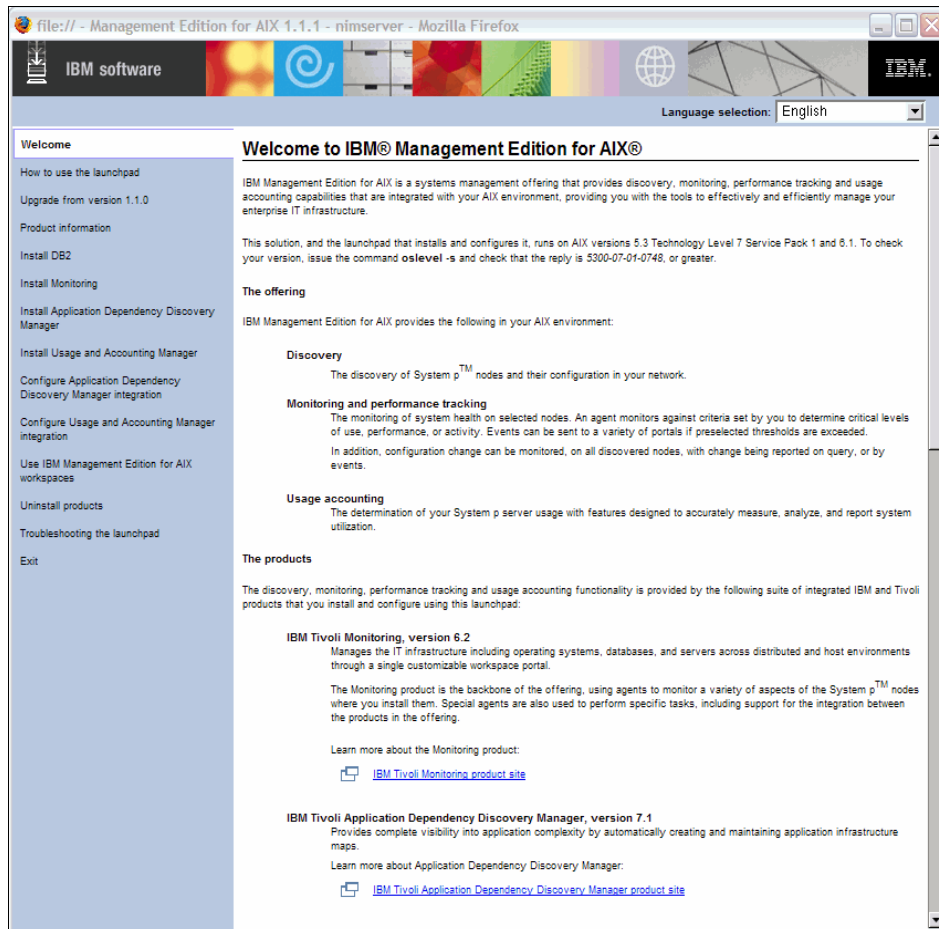
Figure 3-19   Launchpad Initial window

3.3.3 Installing the Tivoli Enterprise Monitoring Server

Figure 3-20 on page 54 shows the Launchpad initial options for installing the IBM Tivoli Monitoring components, as well as an overview of their functions and the sequence in which you must install the individual components. It also has links to more detailed information about using the Launchpad and the IBM Tivoli Monitoring product in the Tivoli Information Center.
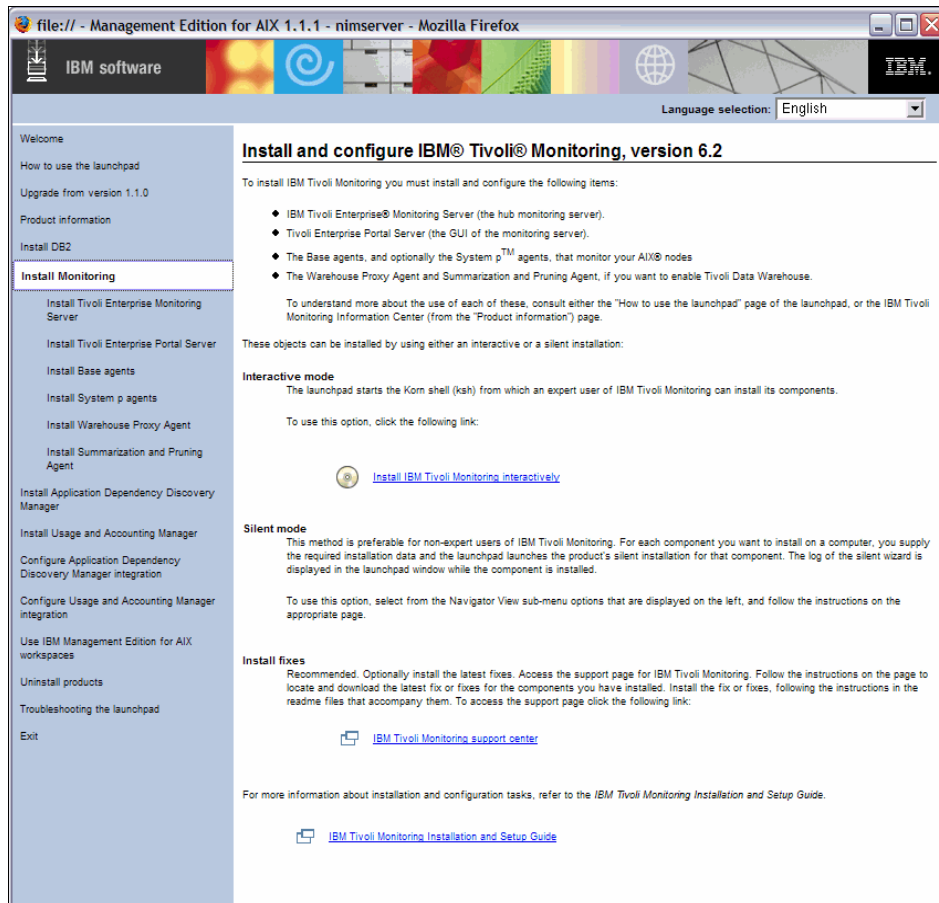
Figure 3-20   Installing Tivoli Monitoring options from the Launchpad

To start the installation:

1. Accept the license agreement if prompted.

   After you accept the licensing agreement, you will be presented with the main Install Monitoring window. From here, the steps that are required to perform a silent installation are laid out logically in the left sidebar. Alternately, you can choose the option "Install IBM Tivoli Monitoring interactively." This option launches a Korn shell (ksh) that prompts you for the required configuration parameters for the installation. This option is essentially identical to installing the product from the command line.

> **Note:** We recommend the Silent Mode installation if you are inexperienced with the IBM Tivoli Monitoring product. This method prompts for the information that the wizard requires to complete an install of the selected component.

Figure 3-21 shows the options for installing and configuring the Tivoli Enterprise Monitoring Server (TEMS) from the Launchpad. From this window, you can either select the components that you want to install or configure from the list in the left sidebar.

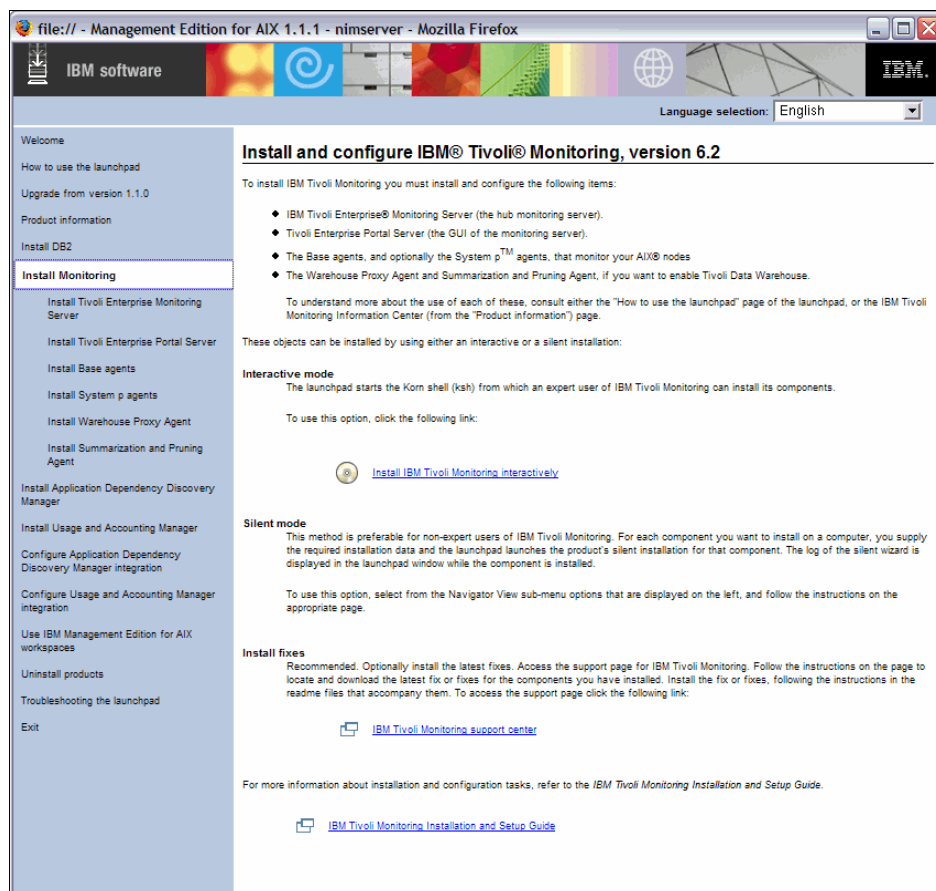The following sections guide you through each of these tasks.

Figure 3-21   Installing the Tivoli Enterprise Monitoring Server

Step 1: Install Tivoli Enterprise Monitoring Server (TEMS)

Clicking "Install Tivoli Enterprise Monitoring Server" in the sidebar launches the window that is seen in Figure 3-22. Here, you are prompted for the information that the wizard requires to proceed with installing the TEMS and the sequence to follow to install the TEMS, to configure the TEMS, and to enable application support for the monitoring agents. Finally, you will be asked to start the TEMS. The required fields are marked with a red asterisk (*):

▶ The install directory (the default is `/opt/IBM/ITM`) is a required field.

▶ The Tivoli Enterprise Monitoring Server name is a required field.

▶ Optionally, you can enter a 32-character encryption key or accept the default.
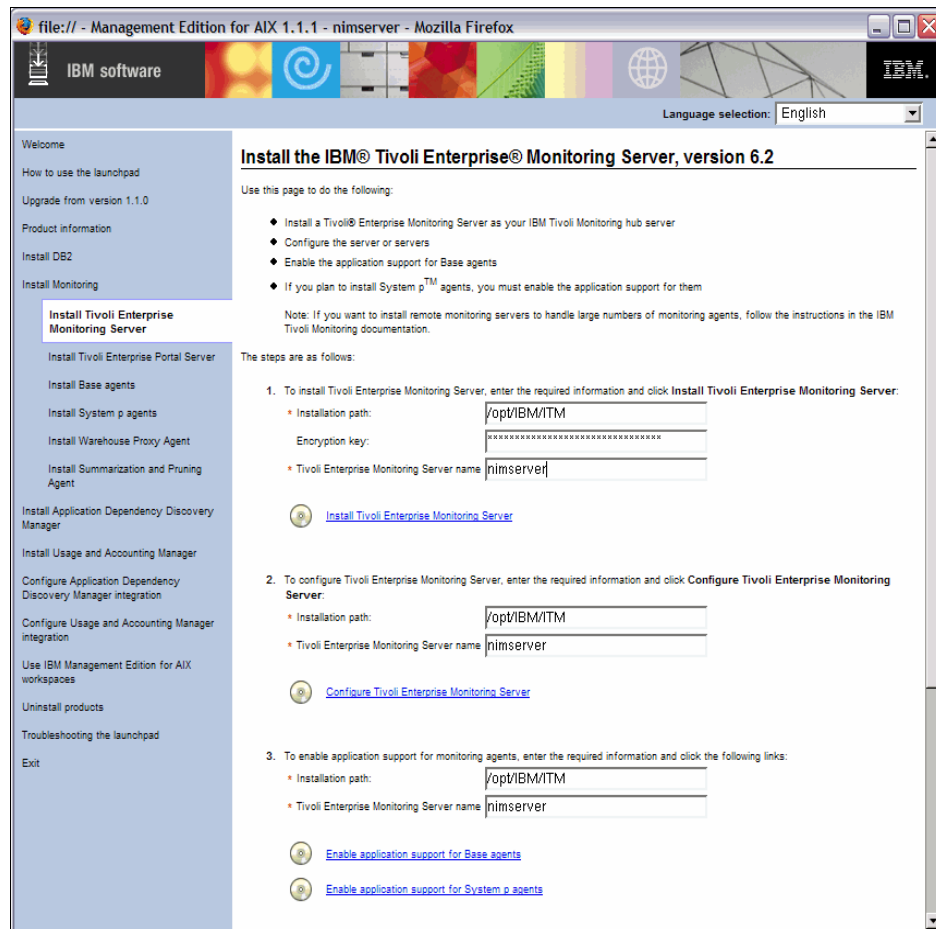

Figure 3-22   Installing the Tivoli Enterprise Monitoring Server

After the install process has initiated, the install log is logged to the display window within the Launchpad while the component is installed. The install process will also identify any of the particular install logs that you need to inspect or review after the component is installed. You can view the entire log by utilizing the scroll bar as shown in Figure 3-23.
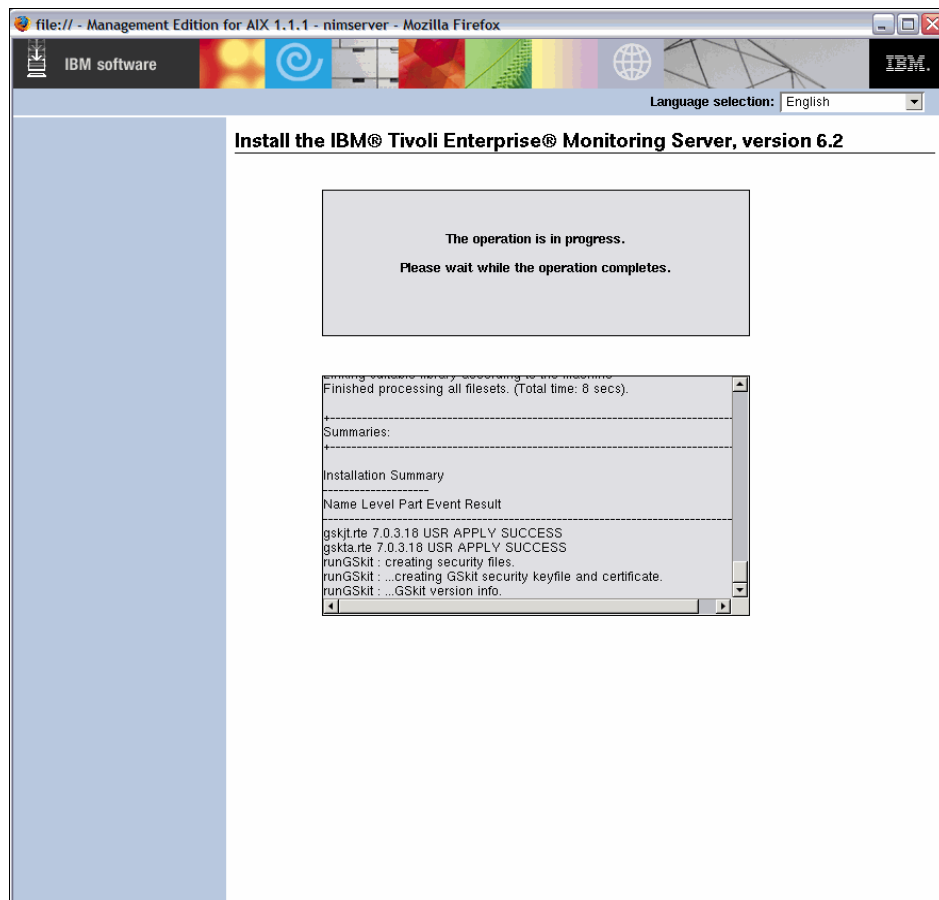

Figure 3-23   Installation progress information

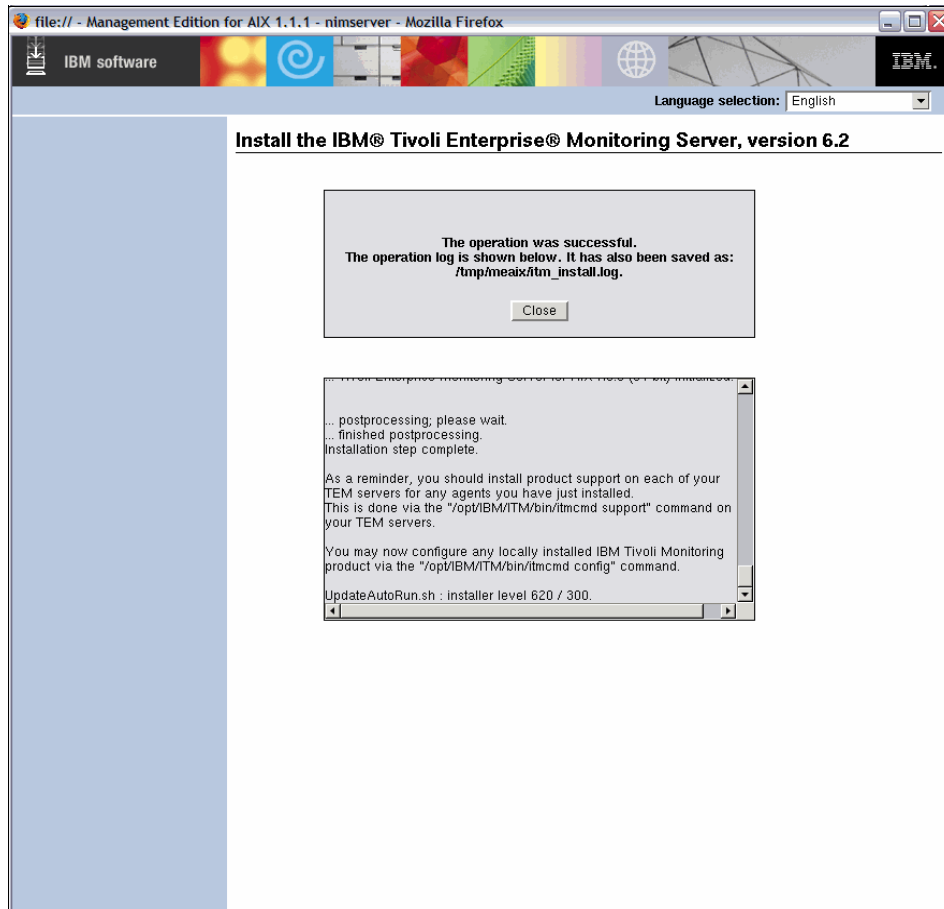The installation progresses, as shown in Figure 3-24.

Figure 3-24   Tivoli Enterprise Monitoring Server Installation

After this installation completes successfully, click the link in Step 2 (Figure 3-25 on page 59).

Step 2: Configuring the Tivoli Enterprise Monitoring Server

Enter the required fields and select **Configure Tivoli Enterprise Monitoring Server** as shown in Figure 3-25. This selection initiates the configuration and prompts you for further information and settings that are required to complete the configuration the TEMS hub.
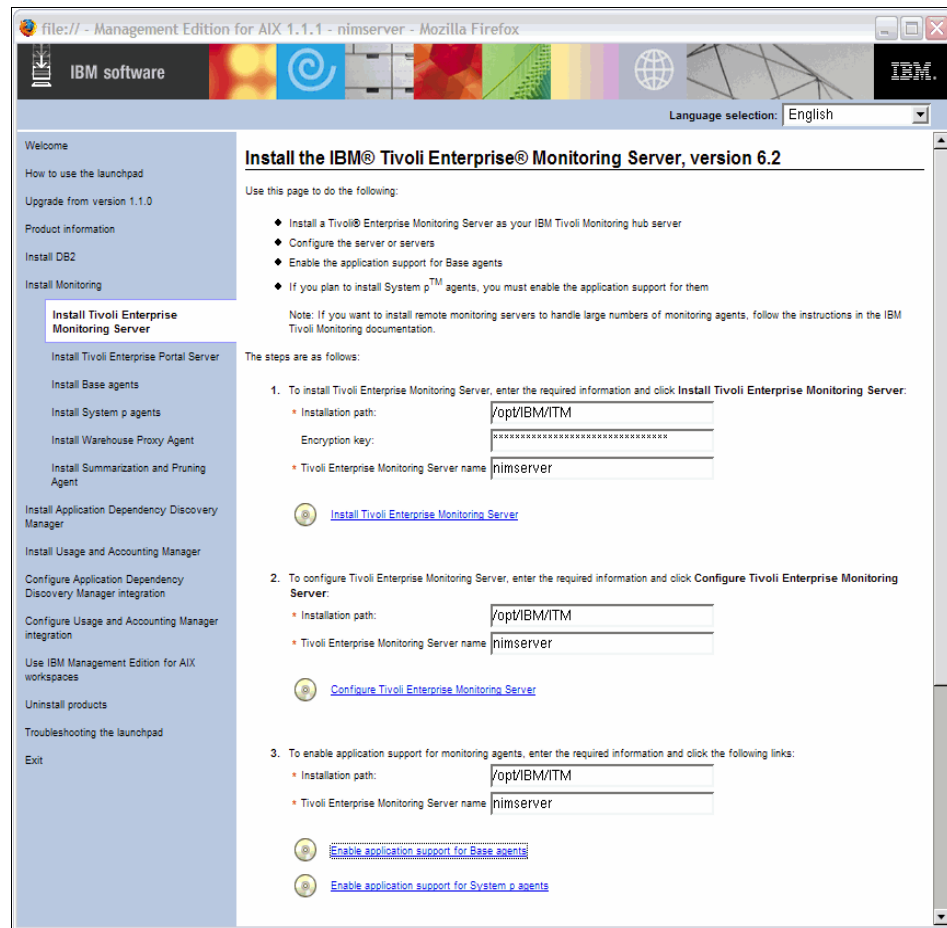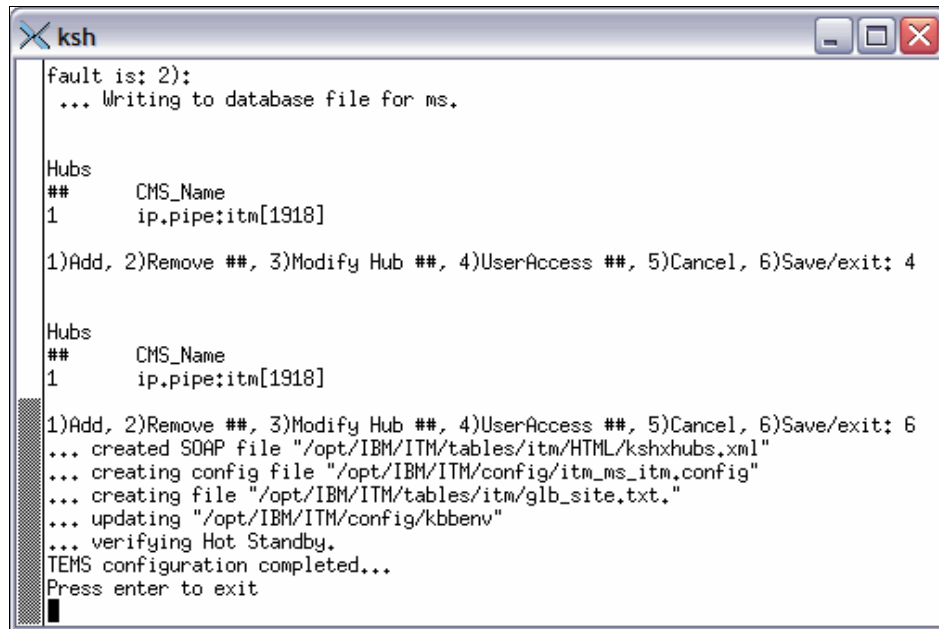

Figure 3-25   Tivoli Enterprise Monitoring Server Installation completion

Next, you enter additional install parameters that are required to finalize the configuration of the TEMS hub as shown in Figure 3-26.


```
fault is: 2):
 ... Writing to database file for ms.


Hubs
##      CMS_Name
1       ip.pipe:itm[1918]

1)Add, 2)Remove ##, 3)Modify Hub ##, 4)UserAccess ##, 5)Cancel, 6)Save/exit: 4


Hubs
##      CMS_Name
1       ip.pipe:itm[1918]

1)Add, 2)Remove ##, 3)Modify Hub ##, 4)UserAccess ##, 5)Cancel, 6)Save/exit: 6
... created SOAP file "/opt/IBM/ITM/tables/itm/HTML/kshxhubs.xml"
... creating config file "/opt/IBM/ITM/config/itm_ms_itm.config"
... creating file "/opt/IBM/ITM/tables/itm/glb_site.txt."
... updating "/opt/IBM/ITM/config/kbbenv"
... verifying Hot Standby.
TEMS configuration completed...
Press enter to exit
```

Figure 3-26   Additional configuration options

Configuration options

The configuration options provided in the Korn shell (ksh) that is spawned by the Launchpad are identical to the configuration options that are displayed if the install is initiated from the command-line interface (CLI) using the install.sh script from the directory that contains the IBM Tivoli Monitoring install images (provided in Example 3-5 on page 61):

1. When prompted for the IBM Tivoli Monitoring V6.2 home directory, press Enter to accept the default (`/opt/IBM/ITM`).

2. If you want to use another installation directory, type the full path to that directory and press Enter.

3. The output shown in Example 3-5 on page 61 is displayed. Type `1` to start the installation and press Enter. The user license agreement is displayed. Press Enter to read through the agreement.

Example 3-5   Tivoli Enterprise Manager Installation Output

```
Select one of the following:

1) Install products to the local host.
2) Install products to depot for remote deployment (requires TEMS).
3) Install TEMS support for remote seeding
4) Exit install.

Please enter a valid number: 1

Initializing ...
install.sh warning: unarchive of "/appo/itm62/unix/jraix526.tar" may
have failed, continuing ...
International Program License Agreement

Part 1 - General Terms
BY DOWNLOADING, INSTALLING, COPYING, ACCESSING, OR USING
THE PROGRAM YOU AGREE TO THE TERMS OF THIS AGREEMENT. IF
YOU ARE ACCEPTING THESE TERMS ON BEHALF OF ANOTHER PERSON
OR A COMPANY OR OTHER LEGAL ENTITY, YOU REPRESENT AND
WARRANT THAT YOU HAVE FULL AUTHORITY TO BIND THAT PERSON,
COMPANY, OR LEGAL ENTITY TO THESE TERMS. IF YOU DO NOT
AGREE TO THESE TERMS,
- DO NOT DOWNLOAD, INSTALL, COPY, ACCESS, OR USE THE
PROGRAM; AND
Press Enter to continue viewing the license agreement, or
enter "1" to accept the agreement, "2" to decline it, "3"
to print it, "4" to read non-IBM terms, or "99" to go back
to the previous screen.
```

4. Type 1 and press Enter to accept the agreement.

5. Press Enter to accept the default encryption key, or type a 32-character encryption key and press Enter.

> **Note:** Follow these tips to select an encryption key:
>
> ► Do not use the equal sign (=), the comma (,), or the l character in the encryption key.
>
> ► Ensure that if you select not to use the default encryption key that you record the key value. This same key will be required during the installation of components that communicate with this monitoring server.

6. Type the number for the operating system on which you install the monitoring server. The default value is your current operating system.

   Press Enter.

7. Type 1 to confirm the operating system and press Enter.

   A numbered list of available components is displayed.

8. Select the number which corresponds to the option for the Tivoli Enterprise Monitoring Server and press Enter.

9. Type 1 to confirm the selection and press Enter.

> **Note:** If you are installing on AIX Version 6.1 and if you only see AIX 5.3 in the selection list, select AIX 5.3. After the closest level to your level is selected, the component will install and operate normally.

At this point, the installation begins as shown in Example 3-6.

Example 3-6   Installation of the Tivoli Enterprise Monitoring Server

```
Product packages are available in /opt/IBM/ITM

Product packages are available for the following operating systems
and component support categories:

1) AIX R5.2 (32 bit)
2) AIX R5.2 (64 bit)
3) AIX R5.3 (32 bit)
4) AIX R5.3 (64 bit)
5) Tivoli Enterprise Portal Browser Client support
6) Tivoli Enterprise Portal Server support
7) Tivoli Enterprise Monitoring Server support

Type the number for the OS or component support category you want,
or type "q" to quit selection
[ number "4" or "AIX R5.3 (64 bit)" is default ]:

You selected number "4" or "AIX R5.3 (64 bit)"
Is the operating system or component support correct [ 1=Yes, 2=No ;
default is "1" ] ? y

The following products are available for installation:

1) Monitoring Agent for UNIX Logs V06.20.00.00
2) Monitoring Agent for UNIX OS V06.20.00.00
3) Summarization and Pruning Agent V06.20.00.00
```

```
4) Tivoli Enterprise Monitoring Server V06.20.00.00
5) Tivoli Enterprise Portal Server V06.20.00.00
6) Tivoli Enterprise Services User Interface V06.20.00.00
7) Universal Agent V06.20.00.00
8) Warehouse Proxy V06.20.00.00
9) all of the above

Type the numbers for the products you want to install, or type "q"
to quit selection.
If you enter more than one number, separate the numbers by a comma
or a space.

Type your selections here: 4

The following products will be installed:

Tivoli Enterprise Monitoring Server V06.20.00.00

Are your selections correct [ 1=Yes, 2=No ; default is "1" ] 1

... installing "Tivoli Enterprise Monitoring Server V06.20.00.00
for AIX R5.3 (64 bit)"; please wait.

=> installed "Tivoli Enterprise Monitoring Server V06.20.00.00 for
AIX R5.3 (64 bit)".
... Initializing component Tivoli Enterprise Monitoring Server
V06.20.00.00 for AIX R5.3 (64 bit).
```

10.At the prompt, enter a host name for your Monitoring Server and press Enter
   as shown in Example 3-7.

Note: Do not use the fully qualified domain name as the host name for your
monitoring server.

Example 3-7   Installation completion

```
Please enter TEMS name [ TEMS is default ]: nimserver
... creating config file
"/opt/IBM/ITM/config/server2_ms_HUB_TEMS.config"
... creating file "/opt/IBM/ITM/tables/HUB_TEMS/glb_site.txt."
... updating "/opt/IBM/ITM/config/kbbenv"
... verifying Hot Standby.
... Tivoli Enterprise Monitoring Server V06.20.00.00 for AIX R5.3
(64 bit) initialized.
```

```
Do you want to install additional products or product support
packages [ 1=Yes, 2=No ; default is "2" ]
... postprocessing; please wait.
... finished postprocessing.
Installation step complete.

As a reminder, you should install product support on each of your
TEM servers for any agents you have just installed.
This is done via the "/opt/IBM/ITM/bin/itmcmd support" command on
your TEM servers.

You may now configure any locally installed IBM Tivoli Monitoring
product via the "/opt/IBM/ITM/bin/itmcmd config" command.
```

The Tivoli Enterprise Monitoring Server is now installed and configured.

You can confirm that the server is installed and functional by using the `cinfo` command as shown in Example 3-8.

Example 3-8   The cinfo command displaying running processes

```
root@nimserver:/# /opt/IBM/ITM/bin/cinfo -r

*********** Fri Mar  6 11:57:26 CST 2009 ******************
User: root Groups: system bin sys security cron audit lp
Host name : nimserver    Installer Lvl:06.20.01.00
CandleHome: /opt/IBM/ITM
**********************************************************
Host       Prod  PID    Owner  Start     ID         ..Status
nimserver  ms    483604 root   18:06:01  nimserver  ...running
```

After you have successfully installed and configured the Tivoli Enterprise
Monitoring Server hub, you can then proceed with Step 3 where you install the
application support on the TEMS hub for both the base agents and the System p
agents as shown in Figure 3-27 on page 65.

Step 3: Enabling application support for agents

When we installed the IBM Tivoli Monitoring V6.2 components, the application
support was installed for the base agents. The agents that come with the
installation for the System p are located on separate media.

After they are installed, you need to enable the application support for these
agents to function correctly. In addition, you need to configure and start the
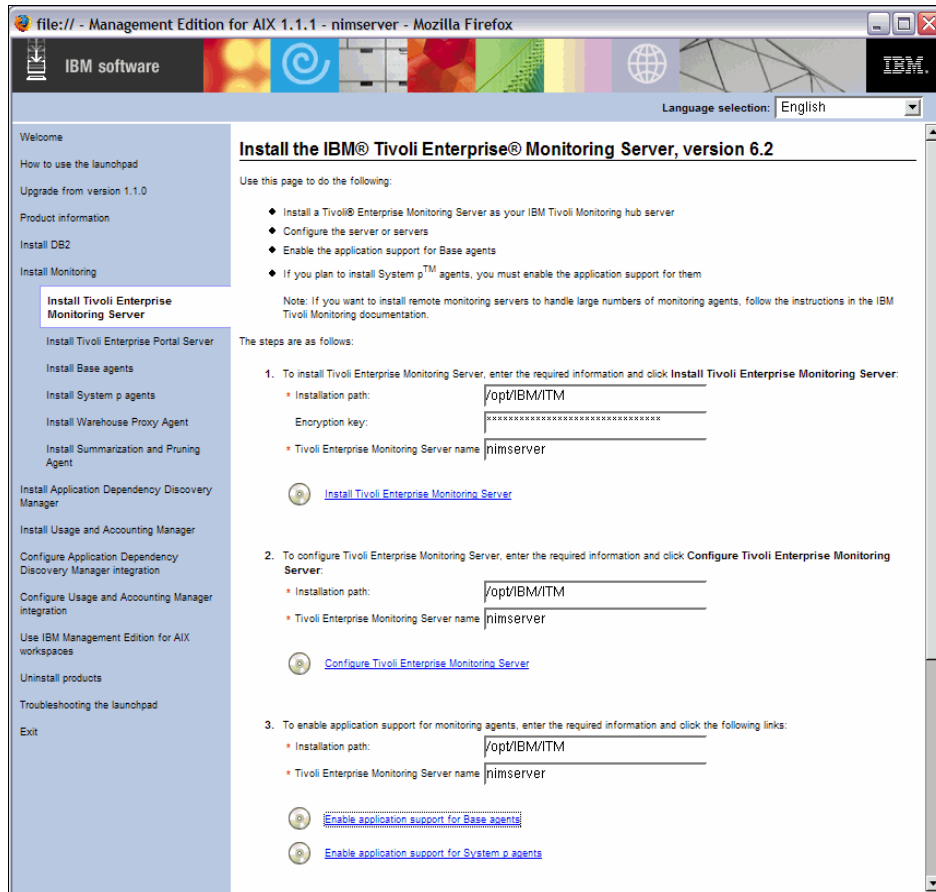agents on the Virtual I/O Server.

Figure 3-27   Installing the application support for the agents

Step 4: Starting the Tivoli Enterprise Monitoring Server

When the TEMS hub is installed and configured and you have installed the application support for the agents, you can now start the TEMS either by clicking the option at Step 4 of the Install Panel, or alternatively, you can the start the TEMS from the command line using the `itmcmd` command.

Example 3-9 on page 66 shows starting and confirming the TEMS from the command line. After the TEMS has been started, you can then confirm that the server process is running by using the `cinfo` command.

Example 3-9   Starting the Tivoli Enterprise Monitoring Server from the command line

```
root@nimserver:/# /opt/IBM/ITM/bin/itmcmd server start nimserver >
/tmp/meaix/startstop_server.log 2>&1
root@nimserver:/# cat /tmp/meaix/startstop_server.log
Starting TEMS...
It can take up to ten minutes.
KCIIN0082E TEMS started...
```

> **Note:** Before proceeding with installing the Tivoli Enterprise Portal, you need to stop the TEMS hub either by using Step 4 on the Tivoli Enterprise Monitoring window or by using the command line and the `itmcmd` command.

3.3.4 Installing the Tivoli Enterprise Portal Server

The second phase of the process is to install and configure the IBM Tivoli Enterprise Portal Server.

Start by selecting **Install Tivoli Enterprise Server** from the Launchpad sidebar. You will see the installation window and the steps, as well as an explanation of the process and the required information and details to launch the installation wizard and configure the required components.

This page assists you to:

► Install the Tivoli Enterprise Portal Server
► Install the Tivoli Enterprise Portal Server support for monitoring agents
► Install the Tivoli Enterprise Portal Server Browser Client support for monitoring agents
► Configure the server

Step 1: Installing the Tivoli Enterprise Server Portal

Use the following steps to install the Tivoli Enterprise Portal Server:

1. Click **Install Tivoli Enterprise Portal Server**.
2. Then, enter the required information in the wizard for the installation directory and the encryption key, or you can use the default encryption key.
3. Click **Install Tivoli Enterprise Portal Server** in the main window as shown in Figure 3-28 on page 67.

> **Note:** The IBM Eclipse Help Server is automatically installed when you install the Tivoli Enterprise Portal.
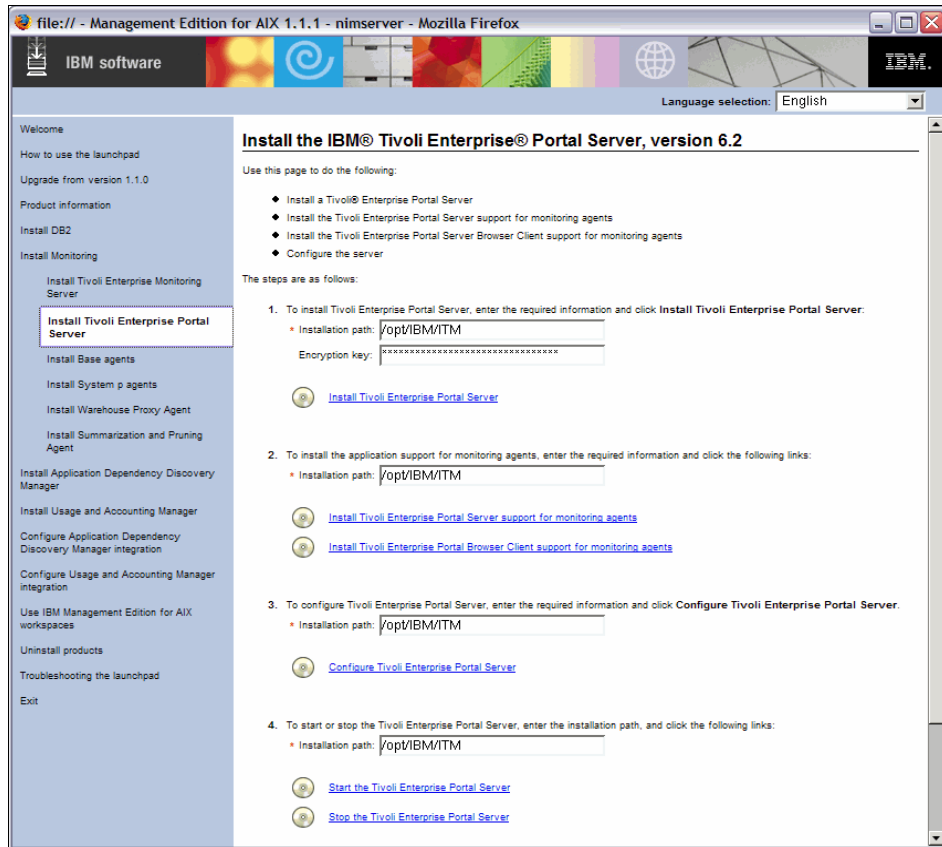
Figure 3-28   Tivoli Enterprise Portal Server installation window

Prior to installation, you are required to accept the license agreement. After you accept the agreement, the wizard prompts you for the installation directory where the IBM Tivoli Monitoring install images are located, which is shown in Figure 3-29 on page 68.

4. Enter the full path to the install images or browse to locate them.
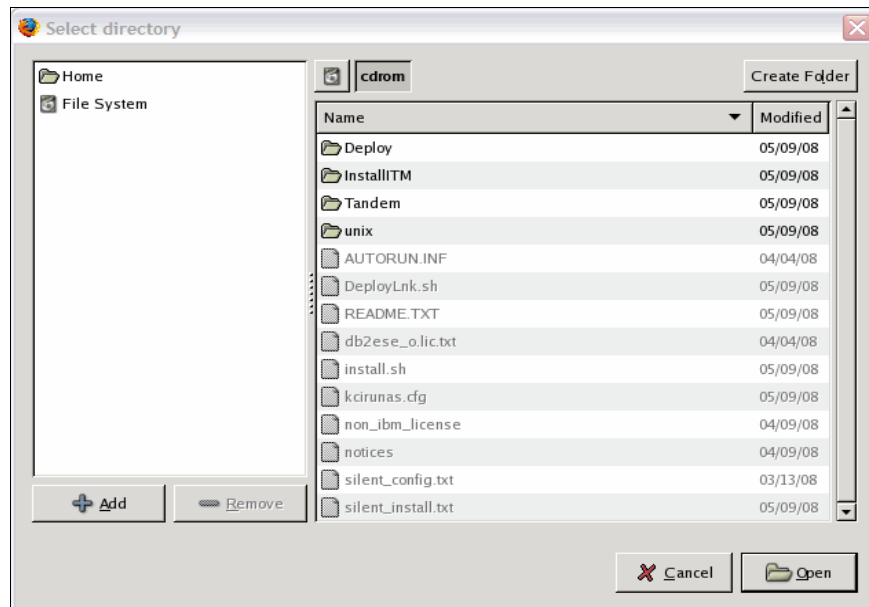
5. Select **Open**.

6. Click **OK** to proceed.

Figure 3-29   Installation setup directory

Next, the wizard installs the Tivoli Enterprise Portal Server and displays the detailed output in the window until the process completes, as shown in Figure 3-30.
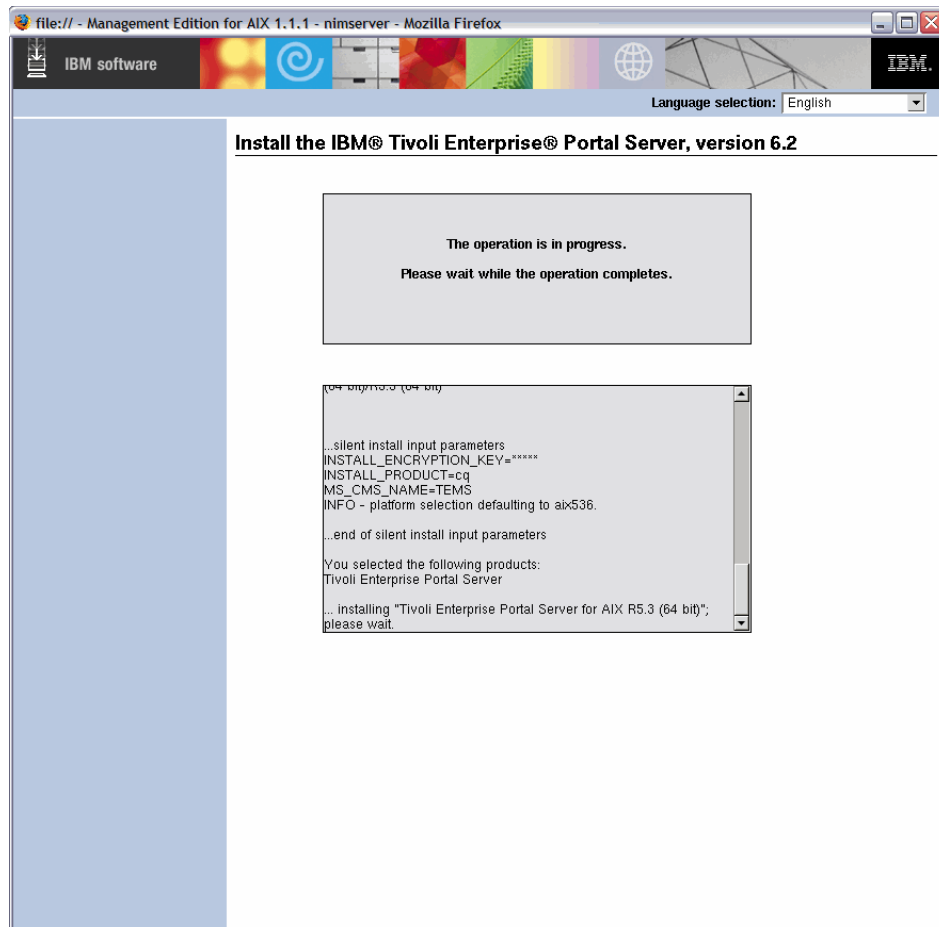
Figure 3-30   Tivoli Enterprise Portal Server installation progress window

In addition to logging in to the window, you are informed about the success or failure of the installation and which logs to review.

Figure 3-31 shows the wizard after the successful installation of the Tivoli Enterprise Portal Server.
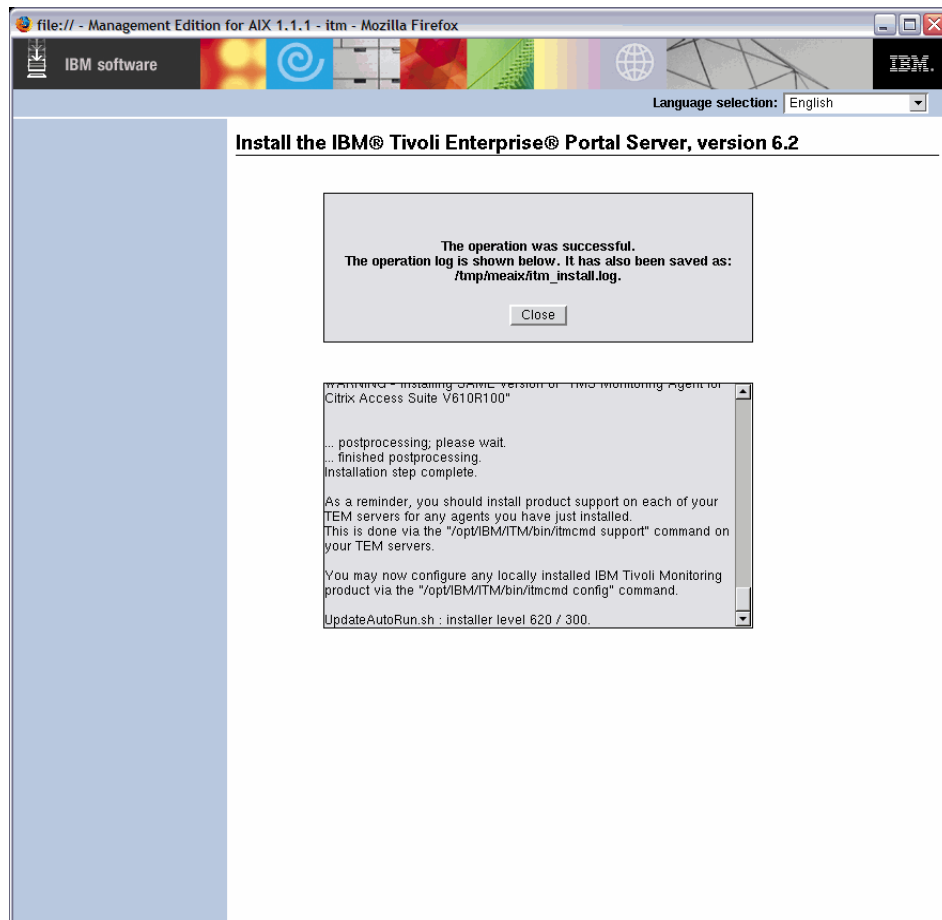
Figure 3-31   Tivoli Enterprise Portal Server installation success window

Step 2: Installing TEPS support for monitoring agents

Next, we guide you through installing the application support that is required for the monitoring agents:

1. After reading and accepting the license agreement, select the installation path for IBM Tivoli Monitoring.

2. Click to install both the Tivoli Enterprise Portal Server support for monitoring agents and the Tivoli Enterprise Portal Server Browser Client support for monitoring agents in that sequence using the Launchpad install wizard window, as shown in Figure 3-32.

Note: The Tivoli Enterprise Portal Browser Client support package is portal server code that supports the browser clients. You must install the browser client support package on the computer where you install the portal server if you want to connect to it using a browser client.
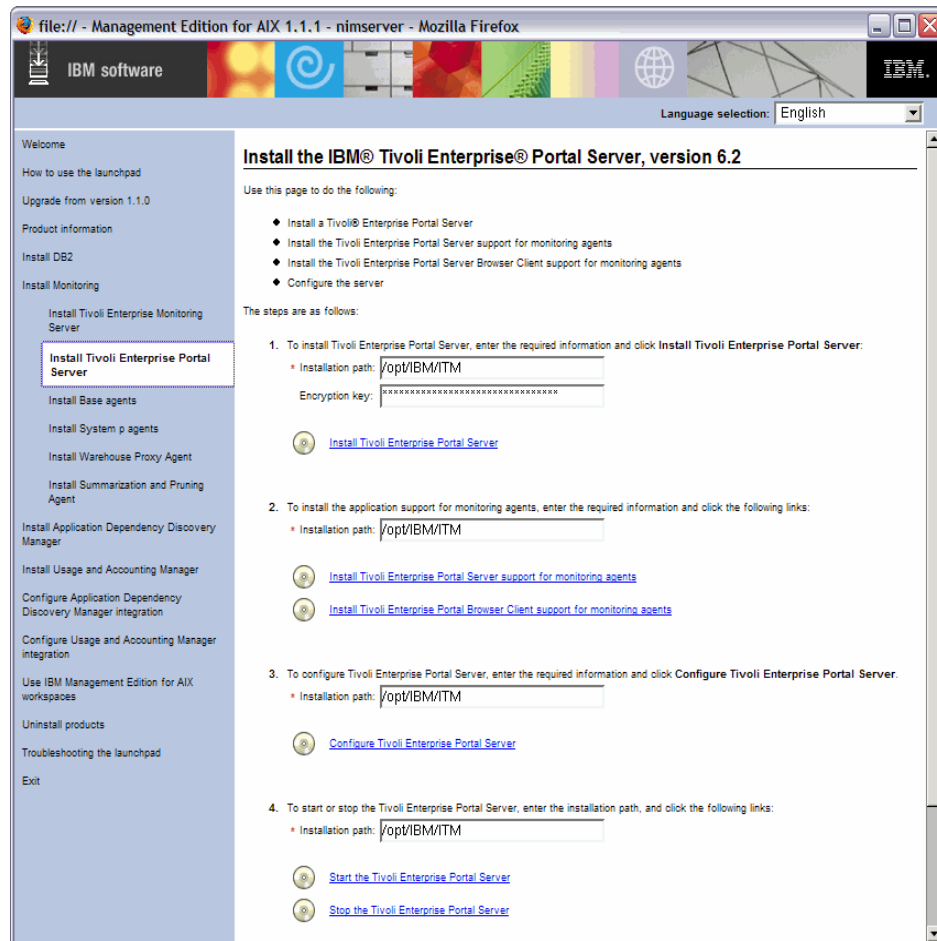
Figure 3-32   Installing the TEPS support for monitoring agents

3. Enter the path for the IBM Tivoli Monitoring setup directory in the dialog box or browse for the location of the setup directory where the install images are, as shown in Figure 3-33.

4. Select **Open** and then click **OK**.

Figure 3-33   Select installation path

Step 3: Configure the Tivoli Enterprise Portal Server

The wizard will now launch a Korn shell (ksh) as shown in Figure 3-34 on page 75, which prompts you for further detailed options to configure the connector settings for the Tivoli Enterprise Portal Server. After this configuration completes, the wizard starts the Install Tivoli Enterprise Portal Server step and connects to the DB2 database instance.

If you are extremely familiar with IBM Tivoli Monitoring and the Tivoli Enterprise Portal Server, you can configure the connector settings for the Tivoli Enterprise Portal Server from the command line using the command `itmcmd config -A cq`.

> **Note:** This part of the process requires you to enter the DB2 and IBM Tivoli Monitoring UIDs, passwords, and the 32-character encryption key in order for the Install Tivoli Enterprise Portal Server to connect to the DB2 instance.

Configuring the Portal Server Connector settings

In the Korn shell (ksh), perform the following steps:

1. Press Enter to accept the default for the following prompt:

   ```
   Edit 'ITM Connector' settings? [ 1=Yes, 2=No ] (default is: 1):
   ```

2. Press Enter to enable the connector.

3. Press Enter to accept the default name of `ITM1` or type your preferred name and press Enter. This name displays in the Common Event Console for this connector.

4. Press Enter to accept the default number of events (`100`) that are to be available in the Common Event Console for this connector, or type the number of events that you want to see displayed and press Enter.

5. Type `2` and press Enter to display only active events in the Common Event Console for this connector. Type `1` and press Enter to view both active and closed events. By default, only active events are displayed.

6. Type `2` and press Enter to skip defining data for extra columns in the Common Event Console.

   When you define a Tivoli Enterprise Console® or Tivoli Netcool/OMNIbus connector, you can define the information that is to be mapped to each of these customizable columns. Refer to the *IBM Tivoli Monitoring Version 6.2.0 Administrator's Guide,* SC32-9408, for information about configuring these connectors.

7. Press Enter when you are asked if the agent connects to a Tivoli Enterprise Monitoring Server.

> **Note:** Although the prompt refers to an agent (which the **itmcmd** command does also), this command is used to configure the portal server.

Configure the connection between the portal server and the hub monitoring server:

1. Type the host name for the hub monitoring server and press Enter.

2. Type the protocol that the hub monitoring server uses to communicate with the portal server. You have four choices: IP.UDP, SNA, IP.PIPE, or IP.SPIPE.

3. If you want to set up a backup protocol, enter that protocol and press Enter. If If you do not want to use a backup protocol, press Enter without specifying a protocol.

4. Depending on the type of protocol that you specified, provide information according to the information shown in Table 2-4 on page 18.

5. Press Enter when you are asked if you want to configure the connection to a secondary monitoring server. The default value is none.

6. Press Enter to accept the default for the Optional Primary Network Name (none).

7. Press Enter to accept the default setting for Secure Sockets Layer (SSL) between the portal server and clients (N).

> **Note:** By default, SSL is disabled. If you want to enable it, type 1 and press Enter.
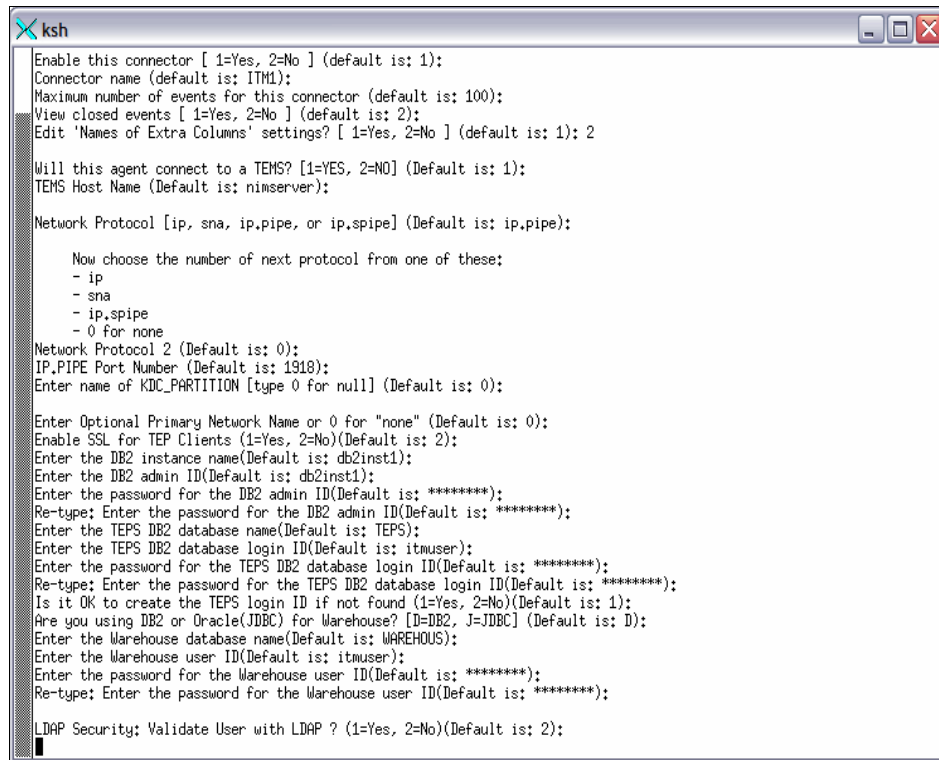
Configure the connection between the portal server and the portal server database.

8. Type the DB2 instance name. The default value is db2inst1. Press Enter.

9. Type the DB2 administrator ID. The default is db2inst1. Press Enter.

> **Note:** The DB2 Administrator account was created during DB2 installation.

10. Type the password for the DB2 administrator ID and press Enter.

11. Confirm the password for the DB2 administrator ID by typing it again. Press Enter.

12. Type the name of the portal server database. The default is TEPS. Press Enter.

13. Type the login name of the database user that the portal server will use to access the database. The default is itmuser. Press Enter.

14. Type the password for the database user and press Enter.

15. Confirm the password for the database user by typing it again. Press Enter.

16. You are now asked if it is okay to create the itmuser login ID if it does not exist. Type 1 and press Enter.

17. You are asked if you are using DB2 or Oracle® for the Tivoli Data Warehouse. Enter D for DB2 or J for Oracle Java Database Connectivity (JDBC). DB2 is the default.

The configuration now proceeds with the parameters that you have entered. You will see a message telling you that InstallPresentation.sh is running and then a message telling you that the installation has completed.


```
Enable this connector [ 1=Yes, 2=No ] (default is: 1):
Connector name (default is: ITM1):
Maximum number of events for this connector (default is: 100):
View closed events [ 1=Yes, 2=No ] (default is: 2):
Edit 'Names of Extra Columns' settings? [ 1=Yes, 2=No ] (default is: 1): 2

Will this agent connect to a TEMS? [1=YES, 2=NO] (Default is: 1):
TEMS Host Name (Default is: nimserver):

Network Protocol [ip, sna, ip.pipe, or ip.spipe] (Default is: ip.pipe):

    Now choose the number of next protocol from one of these:
    - ip
    - sna
    - ip.spipe
    - 0 for none
Network Protocol 2 (Default is: 0):
IP.PIPE Port Number (Default is: 1918):
Enter name of KDC_PARTITION [type 0 for null] (Default is: 0):

Enter Optional Primary Network Name or 0 for "none" (Default is: 0):
Enable SSL for TEP Clients (1=Yes, 2=No)(Default is: 2):
Enter the DB2 instance name(Default is: db2inst1):
Enter the DB2 admin ID(Default is: db2inst1):
Enter the password for the DB2 admin ID(Default is: ********):
Re-type: Enter the password for the DB2 admin ID(Default is: ********):
Enter the TEPS DB2 database name(Default is: TEPS):
Enter the TEPS DB2 database login ID(Default is: itmuser):
Enter the password for the TEPS DB2 database login ID(Default is: ********):
Re-type: Enter the password for the TEPS DB2 database login ID(Default is: ********):
Is it OK to create the TEPS login ID if not found (1=Yes, 2=No)(Default is: 1):
Are you using DB2 or Oracle(JDBC) for Warehouse? [D=DB2, J=JDBC] (Default is: D):
Enter the Warehouse database name(Default is: WAREHOUS):
Enter the Warehouse user ID(Default is: itmuser):
Enter the password for the Warehouse user ID(Default is: ********):
Re-type: Enter the password for the Warehouse user ID(Default is: ********):

LDAP Security: Validate User with LDAP ? (1=Yes, 2=No)(Default is: 2):
```

Figure 3-34   Configuring the Tivoli Enterprise Portal

Step 4: Starting and stopping the Portal Server

You can start or stop the Tivoli Enterprise Portal by clicking the links in Step 4 or by entering the following command from the command line:

> itmcmd agent *<start/stop> <product code>*

The **itmcmd** command can facilitate executing the same command on multiple agents simultaneously from the CLI as shown in Example 3-10.

Example 3-10   The itmcmd multi agent start

```
root@nimserver:/# itmcmd agent start px pk ph cq
Multi-agent start: px
Starting Premium Monitoring Agent for AIX ...
Premium Monitoring Agent for AIX started
```

```
Multi-agent start: pk
Starting Base Monitoring Agent for CEC ...
Base Monitoring Agent for CEC started
Multi-agent start: ph
Starting Base Monitoring Agent for HMC ...
Base Monitoring Agent for HMC started
Multi-agent start: cq
Starting Tivoli Enterprise Portal Server ...
Eclipse Help Server is required by Tivoli Enterprise Portal Server
(TEPS) and will be started...
Eclipse Help Server is already started.
Tivoli Enterprise Portal Server started
Multi-agent start completed
```

Note: If you start the Enterprise Portal Server from the command line, it starts the IBM Eclipse Help Server first if it is not already running.

Table 3-3 on page 77 provides a list of the IBM Tivoli Monitoring product codes and a description of the respective functions of a particular agent or server. You can obtain a thorough and comprehensive list on the system in the file `$INSTALLDIR/registry/proddsc.tbl`.

Table 3-3   IBM Tivoli Monitoring product codes

IBM Tivoli Monitoring product code	Description
ms	Tivoli Enterprise Monitoring Server
ax	IBM Tivoli Monitoring Shared Libraries
cq	Tivoli Enterprise Portal Server
p5	Base Monitoring Agent for AIX
px	Premium Monitoring Agent for AIX
pv	Base Monitoring Agent for Virtual I/O Server
va	Premium Monitoring Agent for Virtual I/O Server
ph	Base Monitoring Agent for Hardware Management Console (HMC)
pk	Base Monitoring Agent for Common Event Console (CEC)
cg	IBM Tivoli Monitoring for Cryptographic Coprocessors
kf	IBM Eclipse Help Server
cw	Tivoli Enterprise Portal Browser Client
hd	Warehouse Proxy
it	Tivoli Enterprise Console (TEC) Graphical User Interface (GUI) Integration
iw	IBM Tivoli Enterprise Portal Server Extensions
cj	Tivoli Enterprise Portal Desktop Client

3.3.5 Install and configure the Base agents
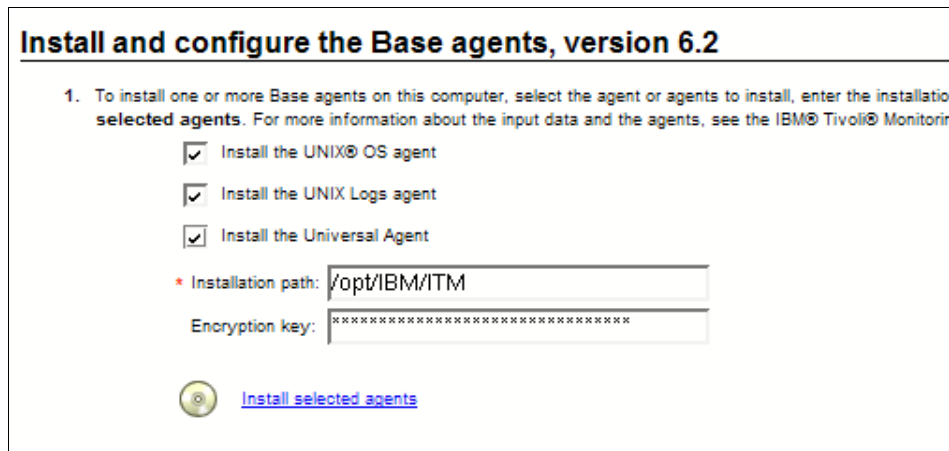
In this next phase, we install the Base agents. You must install and configure the Base agents and System p agents on all of the AIX hosts in the environment that are to be monitored.

Step 1: Installing the Base agents

To install one or more Base agents on this computer:

1. Select the check box for the corresponding agent or agents to install.

2. Enter the installation path and the encryption key.

3. Click **Install selected agents** as shown in Figure 3-35.

Figure 3-35   Installing the Base agents

4. After reading and accepting the licensing agreement, enter the path to the install images or browse for the path to the install images and click **Open**.

5. Then, select **OK** to proceed.

   The install information window will list the agents and their corresponding two character product code as they are installed as shown in Figure 3-36 on page 79.
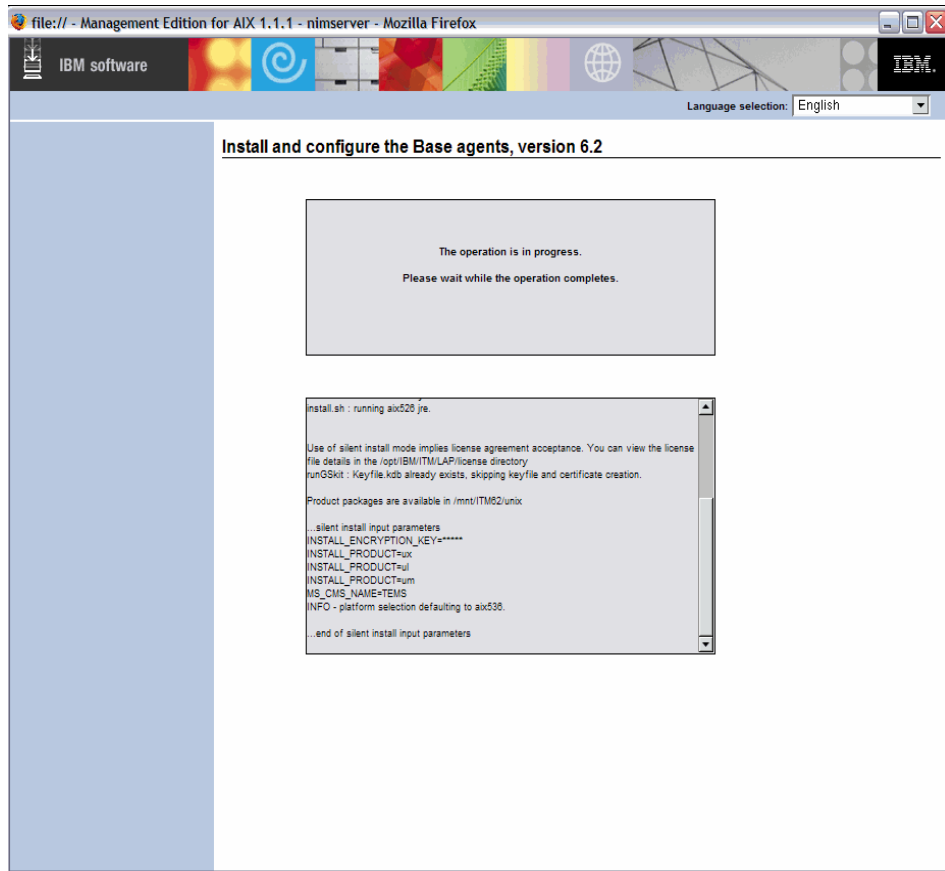
Figure 3-36   Installing Base agents

After the installation completes, the information window appears with information about the success of the installation and where the installation log file can be found as shown in Figure 3-37.
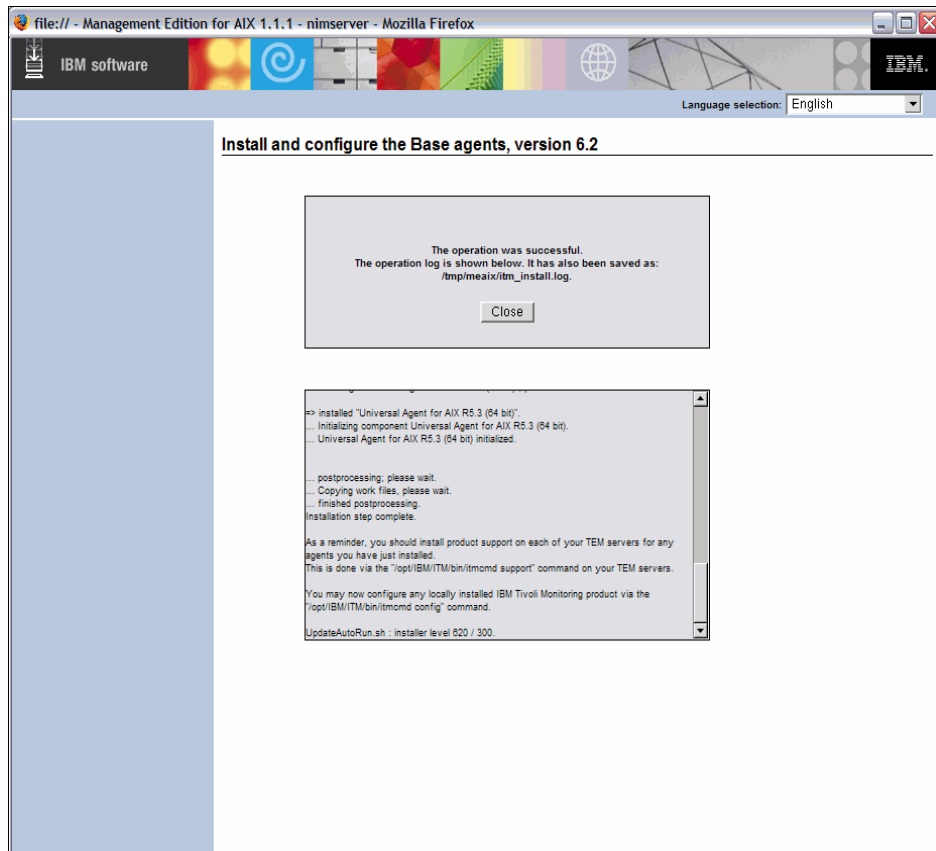
Figure 3-37   Base agents successfully installed

Step 2: Configuring the Base agents

You must now configure each of the Base agents that you have just installed on this host:

1. To configure an agent, enter the installation path of the agent, and click the corresponding link as shown in Figure 3-38.
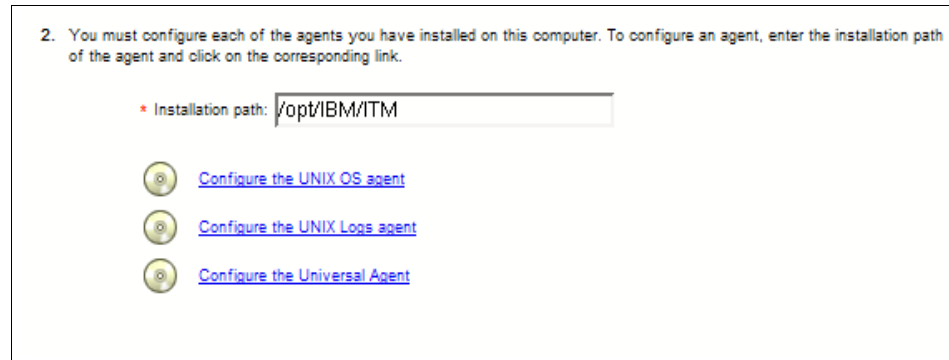
Figure 3-38   Configuration options for the Base agents

Figure 3-39 shows the Korn shell (ksh) agent configuration.
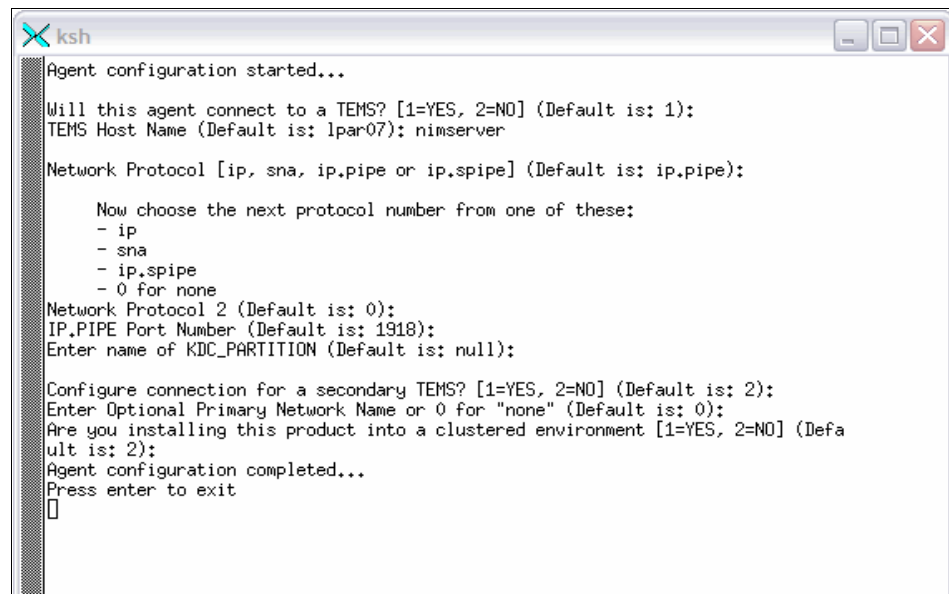
Figure 3-39   Agent configuration shell

You are required to configure the agent parameters individually, because there are slight variations on the parameters that are required for separate agents.

Step 3: Starting the Base agents

You can now start the Base agents by selecting the corresponding check box for the agent and then selecting the link as shown in Figure 3-40.

Figure 3-40   Starting the Base agents

3.3.6 Install and configure System p agents

In this phase, you install the System p agents in the same manner in which you installed the Base agents. You begin by selecting **Install System p agents** in the sidebar of the Launchpad, which will launch the Install window that is shown in Figure 3-41 on page 83.

You must install both the Base agents and System p agents on all AIX hosts to be monitored. In the case of a Workload Partition (WPAR), the agents only need to be installed in the Global Environment. In order to monitor a WPAR setup for application mobility, the agents need to be installed on all systems that can potentially host the WPAR, including the Network File System (NFS) Server.

The systems administrator must be sure to not move the logical partition (LPAR) that has the common event console (CEC) agent installed, because the CEC agent must be on an LPAR of the system that it is monitoring.

Important: Prior to configuring the Hardware Management Console (HMC) Base agent, edit the `$CANDLEHOME/configph_dd.properties` file and add the following lines to the end of the file:

`KQZ_PRODDESC=IBM Tivoli Monitoring: HMC Base Agent`

`KQZ_PRODNAME=IBM Tivoli Monitoring: HMC Base Agent`

Failure to add these lines will generate Java exception errors when you try to configure the HMC Base agent.
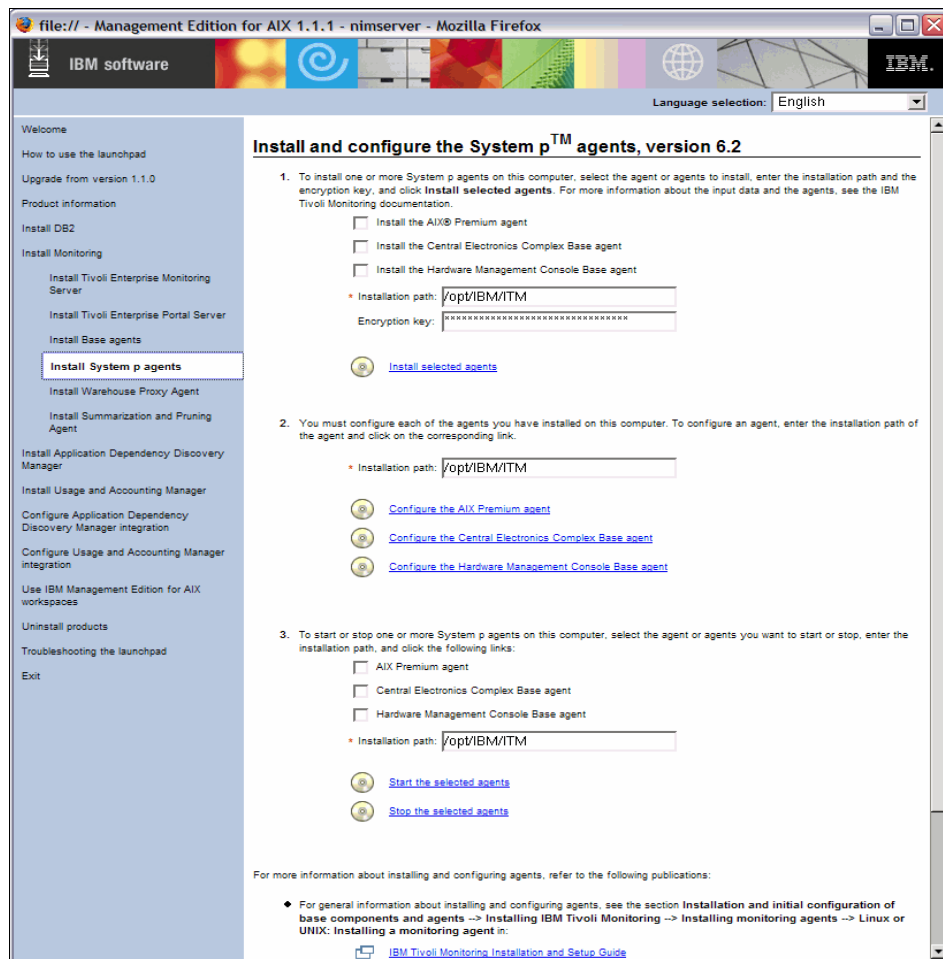
Figure 3-41   Installation options window for System p agents

Step 1: Installing the System p agents

Select the corresponding System p agents to install on the system, and click
Install selected agents to launch the install process as shown in Figure 3-42.
During the process, the information window shows the products or agents being
installed and their corresponding product codes that identify the individual
agents.
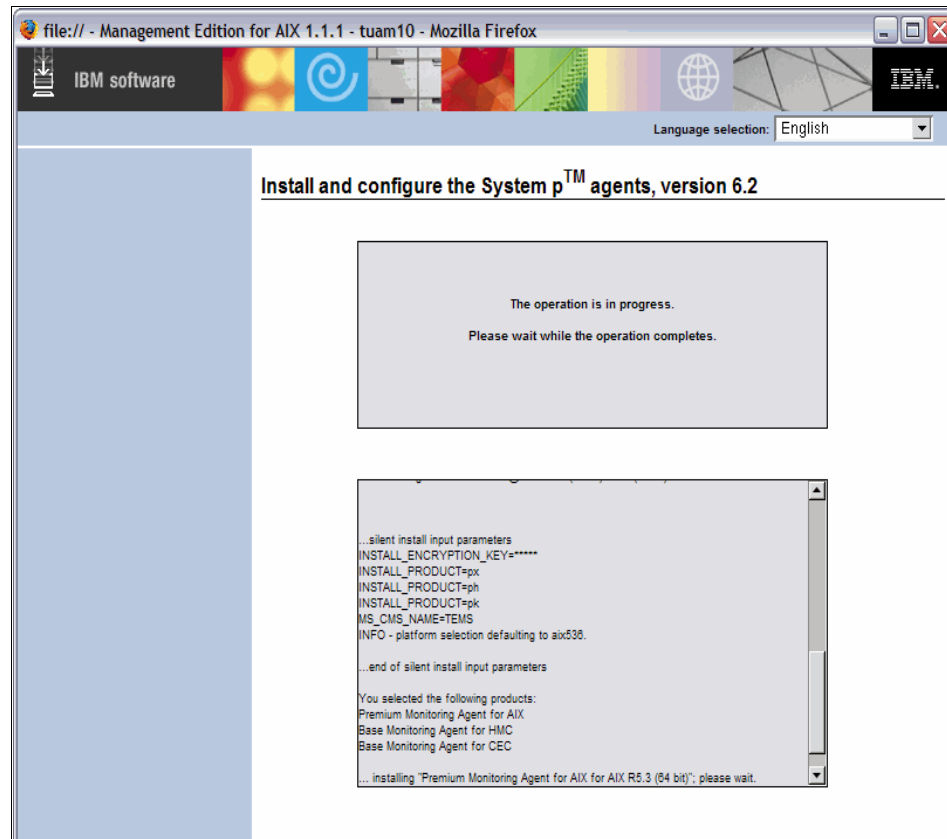
Figure 3-42   System p agents installation

Step 2: Configuring the System p agents

After the installation has completed, configure each agent individually in the "Install and configure the System p agents, version 6.2" window as shown in Figure 3-43.

Figure 3-43   Configure System p agents window

Or, click **Configure <*agent type*>** to launch a Korn shell (ksh), which prompts you for the configuration parameters that are required for the customization of each of the agents as shown in Figure 3-44.

Figure 3-44   Installation parameters for System p agents

Step 3: Starting the agents

You can now select the check box corresponding to the System p agent or agents that are required to start as shown in the Launchpad window in Figure 3-45 on page 87.

Figure 3-45   Starting the System p agents

3.3.7 Installing the Warehouse Proxy agent

To install the Warehouse Proxy agent, enter the required information with the installation path and encryption key as shown in Figure 3-46 on page 88, and then click **Install the Warehouse Proxy Agent**.
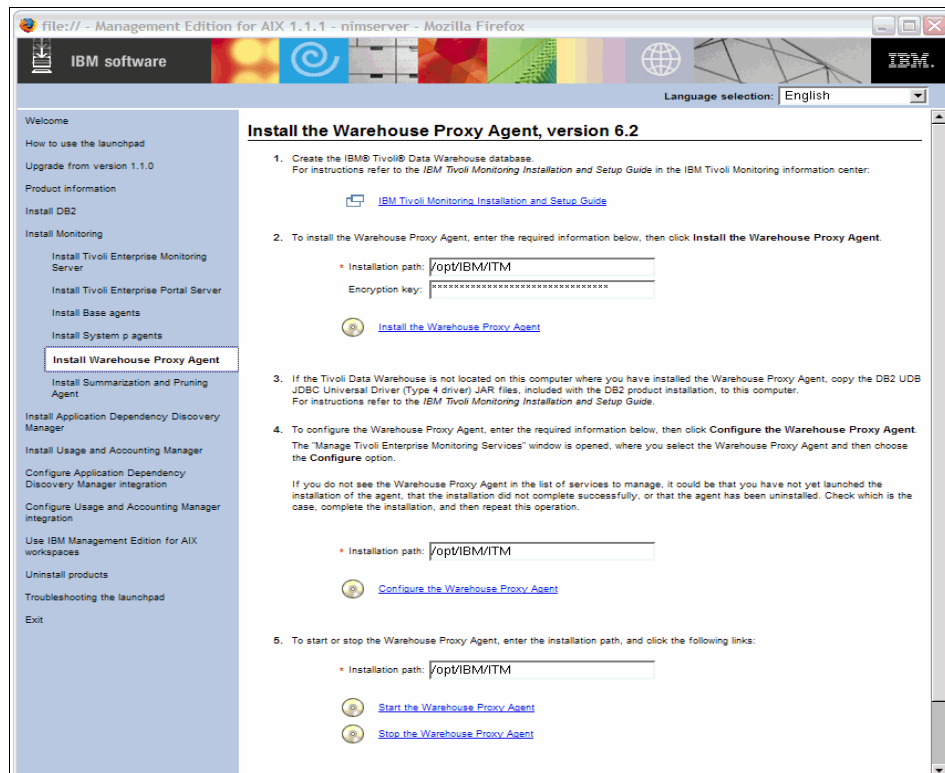
Figure 3-46   Installing the Warehouse Proxy agent

Review and accept the licensing agreement. Then, enter the path to the install images. The window will switch to the install information window as shown in Figure 3-47 until the installation completes.
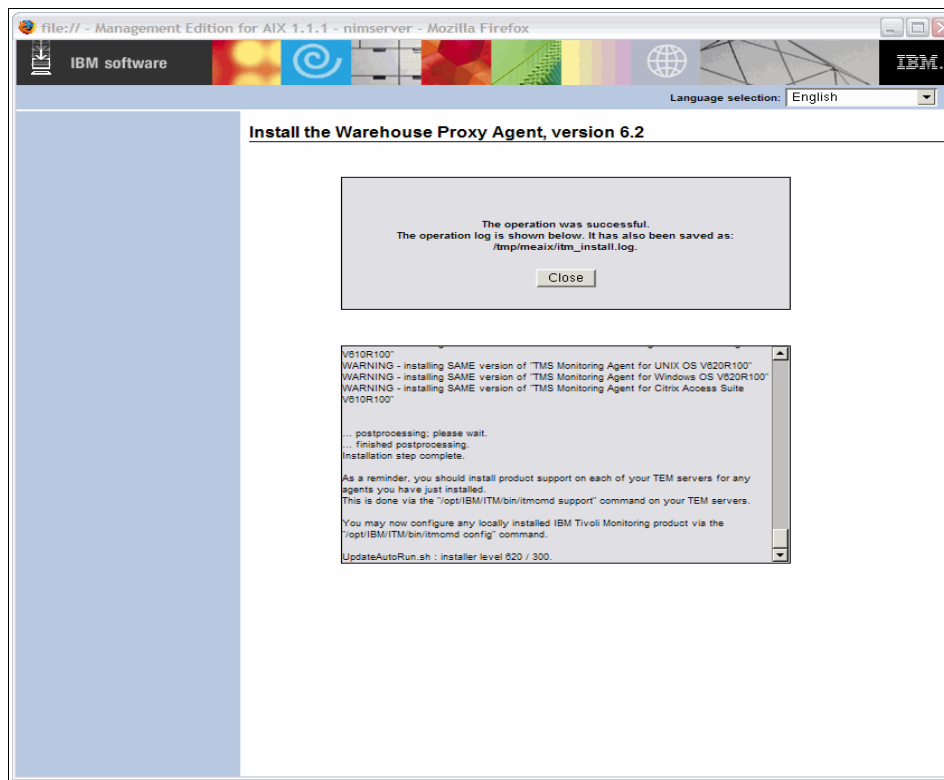
Figure 3-47   Installation information window

To configure the Warehouse Proxy agent, enter the installation directory in Step 3 of the Install Warehouse Proxy agent, and then click **Configure the Warehouse Proxy Agent**.

The Manage Tivoli Enterprise Monitoring Services window opens, where you select **Warehouse Proxy** and then choose **Actions** → **Configure** as shown in Figure 3-48.

> **Note:** If you do not see the Warehouse Proxy agent in the list of services to manage, you might not have launched the installation of the agent, the installation might not have completed successfully, or the agent has been uninstalled. Determine which situation happened, complete the installation, and then repeat this operation.
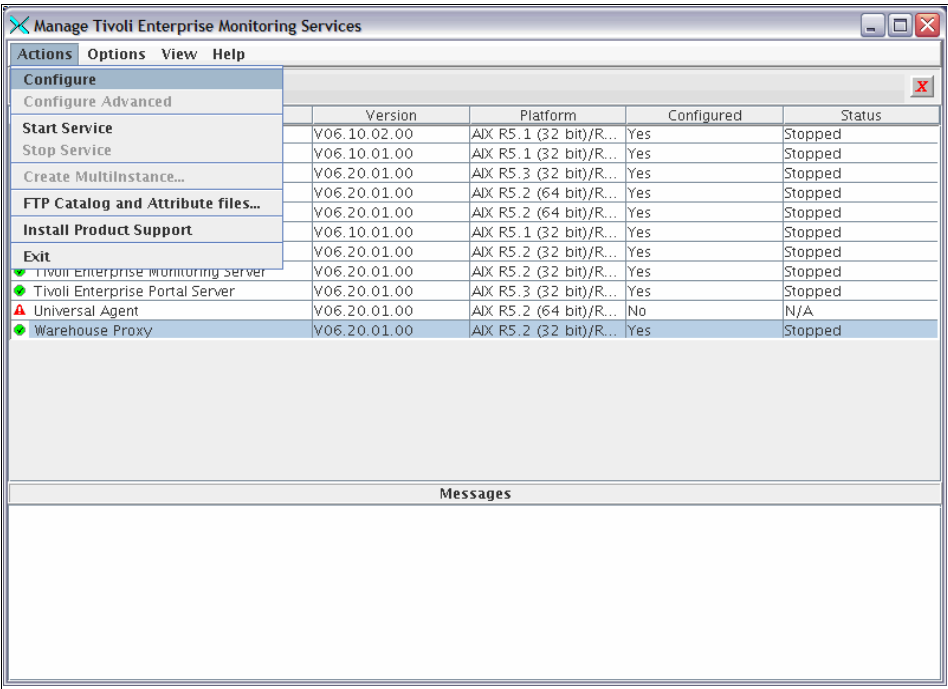
Figure 3-48   Configuring the Warehouse Proxy agent

3.3.8 Installing the Summarization and Pruning agent

To install the IBM Tivoli Monitoring Summarization and Pruning agent, enter the installation directory and encryption key, and then click **Install the Summarization and Pruning Agent** in Step 1 on the LaunchPad window as shown in Figure 3-49.
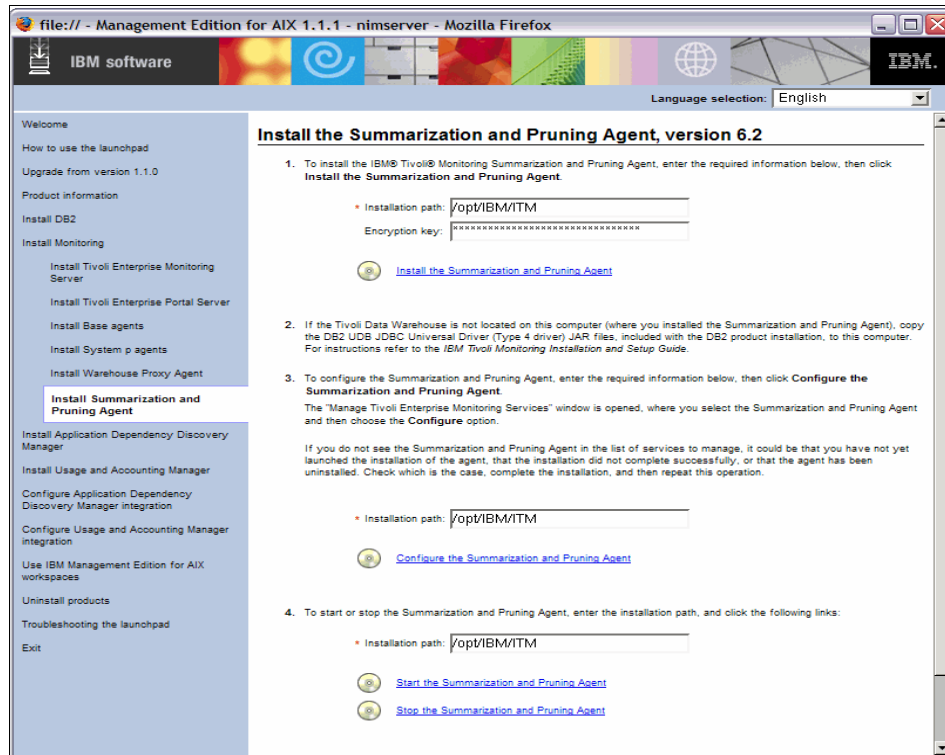
Figure 3-49   Installing the Summarization and Pruning agent

> **Note:** Further information regarding the complex deployments of agents, including remote deployment, is described in detail in Chapter 7, "Deploying IBM Tivoli Monitoring agents in a large scale environment," in *IBM Tivoli Monitoring: Implementation and Performance Optimization for Large Scale Environments*, SG24-7443.

Changing the file permissions

When you install the agents, you install them as the root user, or with privileges greater than those privileges of an applications user.

Follow this procedure to ensure that the permissions are set appropriately:

1. Log in to the host as `root`.

2. Create a user, such as `itm1`, that will own the files and the running agents. Use this command:

   `mkuser itm1`

3. Create a group to which the `itm1` user will belong:

   `mkgroup itmusers`

4. Add the `itm1` UID to the `itmusers` group:

   `chgroup users=itm1 itmusers`

5. List the `itmusers` group to check that the `itm1` UID has been added as shown in Example 3-11.

   Example 3-11   IBM Tivoli Monitoring application Group ID and user ID

   ```
   root@aix:/opt/IBM/ITM# lsgroup itmusers
   itmusers id=201 admin=false users=itm1 registry=files
   ```

6. Run the following command to ensure that the CANDLEHOME environment variable correctly identifies the IBM Tivoli Monitoring V6.2 installation directory:

   `echo $CANDLEHOME`

7. Change to the directory that is returned in the previous step:

   `cd $CANDLEHOME`

8. Confirm that you are in the correct directory:

   `pwd`

9. Run the following commands:

   ```
   chown -R itm1.itmusers .
   chmod -R o-rwx .
   ```

   Important: Running the following steps in the wrong directory can change the permissions on every file in every file system on the host.

10. Run the following command to change the ownership of additional agent files:

    `./bin/SetPerm`

11. The $CANDLEHOME directory now looks Example 3-12.

Example 3-12   User and Group ownership in the $CANDLEHOME directory

```
root@aix:/opt/IBM/ITM# ll
total 136
drwxr-x---   29 itm1     itmusers       4096 Mar 15 10:09 .
drwxr-xr-x    5 bin      bin             256 Mar 10 16:38 ..
drwxrwx---   15 itm1     itmusers       4096 Mar 13 2008  HTML
drwxrwx---    3 itm1     itmusers        256 Mar 13 15:15 InstallITM
drwxr-x---    4 itm1     itmusers        256 Jan 22 2005  JRE
-rw-r--r--    1 itm1     itmusers         86 Mar 15 10:21 KPHOUT.LOG
-rw-r--r--    1 itm1     itmusers         37 Mar 15 10:08 KPKOUT.LOG
-rw-r--r--    1 itm1     itmusers         37 Mar 15 10:09 KPXOUT.LOG
drwxrwx---    3 itm1     itmusers        256 Mar 13 15:24 LAP
drwxrwx---    6 itm1     itmusers        256 Mar 13 15:26 aix513
drwxrwx---    7 itm1     itmusers        256 Mar 13 15:16 aix526
drwxrwx---    3 itm1     itmusers       4096 Mar 13 15:37 bin
drwxr-x---    2 itm1     itmusers       4096 Mar 13 15:12 classes
drwxrwx---    4 itm1     itmusers      12288 Mar 15 10:23 config
drwxr-x---    2 itm1     itmusers        256 Mar 13 15:12 doc
drwxr-x---    2 itm1     itmusers        256 Mar 13 15:12 download
drwxr-x---    6 itm1     itmusers        256 Mar 13 15:24 gskit
drwxr-x---    4 itm1     itmusers       4096 Mar 13 2008  images
drwxr-x---    2 itm1     itmusers        256 Mar 13 15:14 keyfiles
drwxrwx---    4 itm1     itmusers        256 Mar 13 15:25 licenses
drwxrwx---    2 itm1     itmusers       8192 Mar 15 10:11 logs
drwxr-x---    2 itm1     itmusers        256 Mar 13 15:08 lost+found
drwxr-x---    2 itm1     itmusers        256 Mar 13 2008  misc
drwxr-x---    2 itm1     itmusers        256 Mar 13 15:24 old_registry
drwxr-x---    2 itm1     itmusers       4096 Mar 13 15:28 registry
drwxrwx---    2 itm1     itmusers        256 May 10 2008  samples
drwxrwx---    2 itm1     itmusers        256 Mar 13 15:15 sqllib
drwxr-x---    2 itm1     itmusers        256 Mar 13 15:12 tables
drwxrwx---    2 itm1     itmusers        256 Mar 13 15:15 teclib
drwxrwx---    9 itm1     itmusers        256 Mar 13 15:25 tmaitm6
drwxrwx---    2 itm1     itmusers      12288 Mar 15 10:23 tmp
drwxr-x---    2 itm1     itmusers        256 Mar 13 15:12 xfer
```

12. We now restart the agents by either logging in as the `itm1` user or by issuing
    the following command:

    `su - itm1 -c /opt/IBM/ITM/bin/itmcmd agent start all`

13. By using the **cinfo** command, we check to confirm that the agents are
    running and that their processes are owned by our new application user as
    shown in Example 3-13 on page 94.

Example 3-13   Showing itm1 running agents

```
root@aix:/opt/IBM/ITM# ./bin/cinfo -r

*********** Sun Mar 15 10:57:41 CDT 2009 ******************
User: root Groups: system bin sys security cron audit lp
Host name : aix  Installer Lvl:06.20.01.00
CandleHome: /opt/IBM/ITM
**********************************************************
Host  Prod  PID     Owner  Start     ID    ..Status
aix   ph    512006  itm1   10:42:59  None  ...running
```

3.3.9 Configuring and starting the Virtual I/O Server agent

You can now configure the monitoring agent for the Virtual I/O Server:

1. List the available monitoring agents on the Virtual I/O Server:

   ```
   $ lssvc
   ITM_base
   ITM_premium
   TSM_base
   ITUAM_base
   TPC
   DIRECTOR_agent
   ```

2. From the results of the previous command, decide which agent you want to configure:

   ```
   ITM_premium
   ```

3. Run the following command:

   ```
   $ cfgsvc -ls ITM_premium
   MANAGING_SYSTEM
   HOSTNAME
   RESTART_ON_REBOOT
   ```

4. Using the **cfgsvc** command, configure the agent attribute:

   ```
   $ cfgsvc ITM_premium -attr HOSTNAME=nimserver RESTART_ON_REBOOT=TRUE
   MANAGING_SYSTEM=hmc1.itso.ibm.com
   Agent configuration started...
   Agent configuration completed...
   ```

5. Start the monitoring agent using the **startsvc** command:

   ```
   startsvc ITM_premium
   ```

Next, in order for the monitoring agent to gather data from the HMC, you must:

> **Note:** After you configure a Secure Shell connection for one agent, you do not need to configure it for subsequent agents.

1. Determine the name of the managed system that hosts the Virtual I/O Server where the monitoring agent is located.

2. Obtain the public key for the Virtual I/O Server using the following command:

   ```
   viosvrcmd -m <managed system> -p <vios name>-c "cfgsvc -key <ITM
   agent name>"
   ```

   For example:

   ```
   hscroot@hmc1:~> viosvrcmd -m 9117-MMA-SN101F170-L10 -p VIOS1_L10 -c
   "cfgsvc -key ITM_premium"
   ssh-rsa
   AAAAB3NzaC1yc2EAAAABIwAAAQEAnCD5GAC3tKxv7flA4PTrHl5N57ZT+laoKOUVY9rH
   LQQOfPeerykvwJO/YUhg51D4dVD1/uN8lUnvbGrXTjVq4QwEGvOt3kwy9+ACI8E6xp/Q
   HugUoBzSyJMPd3FUPeXHBTvdyYO46wFak8AP7pHjuf+q3xvBOU/e3ax9rmv5oTcaNbhn
   a5mclVuIy1ltbegsduwnvA3udaDpeOev733gY5aTySGGai3BRhImkGFFHbW1lqpO6IG9
   H8N8b1ztOhwY4rgEYE+O9Nizx4xvmz5Z/DycA/xr/e1hhiRFKGpP2pg+RzfatKXIFGVE
   FQg1l7f6iKi55WPJ9cpHtcUyLYsQIQ== root@vios1_l10
   ```

   Where:

   – *Managed system* is the name of the managed system on which the Virtual I/O Server and monitoring agent are located.

   – *vios name* is the name of the Virtual I/O Server LPAR as defined on the HMC.

   – *ITM agent name* is the monitoring agent, for example, `ITM_premium`.

3. Update the authorized_key2 file on the HMC using:

   ```
   mkauthkeys --add public_key
   ```

   Where *public_key* is the output obtained in Step 2. For example:

   ```
   hscroot@hmc1:~> mkauthkeys --add 'ssh-rsa
   AAAAB3NzaC1yc2EAAAABIwAAAQEAnCD5GAC3tKxv7flA4PTrHl5N57ZT+laoKOUVY9rH
   LQQOfPeerykvwJO/YUhg51D4dVD1/uN8lUnvbGrXTjVq4QwEGvOt3kwy9+ACI8E6xp/Q
   HugUoBzSyJMPd3FUPeXHBTvdyYO46wFak8AP7pHjuf+q3xvBOU/e3ax9rmv5oTcaNbhn
   a5mclVuIy1ltbegsduwnvA3udaDpeOev733gY5aTySGGai3BRhImkGFFHbW1lqpO6IG9
   H8N8b1ztOhwY4rgEYE+O9Nizx4xvmz5Z/DycA/xr/e1hhiRFKGpP2pg+RzfatKXIFGVE
   FQg1l7f6iKi55WPJ9cpHtcUyLYsQIQ== root@vios1_l10'
   ```

After you complete these tasks, you can view the data that the monitoring agent has gathered in the Tivoli Enterprise Console.

3.3.10 Viewing the Tivoli Enterprise Console

You can now launch a client browser, Internet Explorer® V5 or higher. In the address field, enter the URL that is shown in Figure 3-50.

Figure 3-50   Tivoli Enterprise Console login window

After you log in, the enterprise environment view workspace appears that is shown in Figure 3-51.
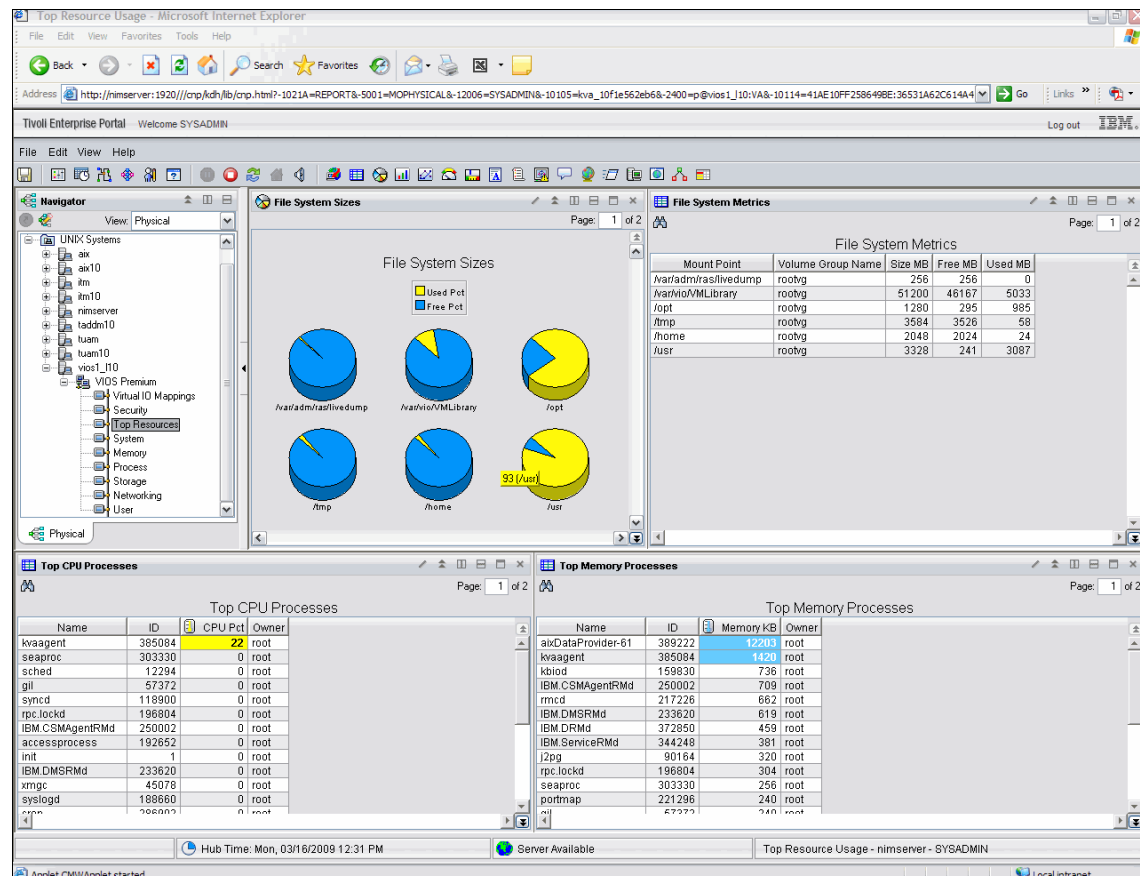
Figure 3-51   Enterprise view of the environment

We describe the Tivoli Enterprise Console and managing the workspaces or views and queries in the console in greater detail in Chapter 4, "Operations" on page 159.

3.4 Installing and configuring TADDM

In this section, we provide steps to install Tivoli Application Development Discovery Manager (TADDM) on AIX:

► 3.4.1, "Preparing AIX" on page 98

3.4.1 Preparing AIX

The test systems referenced in this publication have five logical partitions on two physical servers. One of the partitions is the TADDM server, and the other partitions are target servers. Even though there can be multiple TADDM servers in a large environment, we will have one TADDM server, and it will manage target partitions. If you are interested in configuring multiple TADDM servers for better performance or other management purposes, refer to *Deployment Guide Series: IBM Tivoli Application Dependency Discovery Manager V7.1*, SG24-7616.

Before starting to install DB2 and TADDM, perform the following steps at the operating system level:

1. Install the base operating system (AIX 6.1).

2. Install the Common Desktop Environment (CDE) bundle for X Windows.

3. Install a Web browser (Firefox 2.0) as shown in "Installing Firefox" on page 30

4. Install `unzip` and `sudo` on the TADDM server (After installing `sudo`, type the `visudo` command and then add in the User privilege specification as shown in Example 3-14:

   ```
   cmdbadm ALL=(ALL)        NOPASSWD:ALL
   ```

Example 3-14   Installing and configuring sudo

```
taddm10:/mnt/Linux_tools/ppc ] ls -l sudo*
-rw-r-----   1 root     system        205789 Feb 26 11:00
sudo-1.6.9p15-2.aix5.2.ppc.rpm
-rw-r-----   1 root     system        200394 Feb 26 11:00
sudo-1.6.9p15-2noldap.aix5.2.ppc.rpm

taddm10:/mnt/Linux_tools/ppc ] rpm -ivh
sudo-1.6.9p15-2noldap.aix5.2.ppc.rpm
sudo
#################################################

taddm10:/mnt/Linux_tools/ppc ] visudo
# sudoers file.
#
```

```
# This file MUST be edited with the 'visudo' command as root.
# Failure to use 'visudo' may result in syntax or file permission
errors
# that prevent sudo from running.
#
# See the sudoers man page for the details on how to write a sudoers
file.
## Host alias specification

# User alias specification

# Cmnd alias specification

# Defaults specification

# Runas alias specification

# User privilege specification
root    ALL=(ALL) ALL
cmdbadm ALL=(ALL)        NOPASSWD:ALL

# Uncomment to allow people in group wheel to run all commands
```

5. Install d **ssh** and **lsof** on all the target servers as well as TADDM server. (The ssh needs to be configured on both TADDM and target servers.)

6. Create a file system for TADDM (size: 3 GB): `/opt/IBM/cmdb` on a non-rootvg volume group.

Note: The **unzip** and **sudo** programs are required on the TADDM server, and the **ssh** and **lsof** programs are necessary on all the target servers, as well as the TADDM server. You can find the `ssh` filesets in the Expansion Pack, while the `unzip`, `sudo`, and `lsof` files are in AIX Toolbox for Linux.

3.4.2 Install DB2

The TADDM installation supports several installation scenarios, including:

1. Simple installation with the installation of a DB2 database. This scenario installs and configures a local DB2 database as part of the TADDM installation.

2. Simple installation without the installation of a DB2 database. This scenario is for situations where you already have your local DB2 database installed.

3. Advanced installation with a remote DB2 database. This scenario requires you to install your DB2 database on a remote server prior to the installation of TADDM.

For our implementation, we chose the second option in the previous list of options, which is the simple installation without the installation of a DB2 database. This installation scenario is relatively straightforward, and it can help you understand the relationship between DB2 and TADDM.

Refer to "Install DB2 Enterprise Server" on page 35 for installing a database.

3.4.3 Create DB2 database

In the TADDM DVD, there is a script named `make_db2_db.sh` that helps you to create the initial database for TADDM.

To run **make_db2_db.sh**, log on (or **su**) as the instance owner, and copy the script from $TADDM_MEDIA/TADDM/support/bin or to the home directory of the instance owner of db2inst1. And then, run the **make_db2_db.sh** as shown in Example 3-15. We will create a database that has a name of cmdb.

Example 3-15   Creating database for TADDM

```
$ cp /mnt/TADDM71/TADDM/support/bin/make_db2_db.sh /home/db2inst1
$ ls -l /home/db2inst1
total 24
-rwxr-xr-x    1 db2inst1 db2grp1        4288 Mar  3 10:35 make_db2_db.sh
drwxrwsr-t   16 db2inst1 db2grp1        4096 Mar  3 10:28 sqllib

$ /home/db2inst1/make_db2_db.sh cmdb
```

3.4.4 Install TADDM 7.1

To complete the installation of TADDM:

1. Log in to the AIX system as the root user.

2. Locate the installation media and copy it to the AIX system. We will use the /mnt/TADDM71 directory as a TADDM installation image.

3. Go to the directory where you have the Launchpad directory and run the **launchpad.sh** command, for example:

    # /mnt/launchpad/launchpad.sh

4. Issue the **launchpad.sh** command to display the Launchpad panel, which is shown in Figure 3-52 on page 101.

5. Click **Install Application Dependency Discovery Manager** → **Install product** on the left pane, and then click **Install Tivoli Application Dependency Discovery Manager**.

> **Note:** If you want to invoke an install program for TADDM directly instead of using Launchpad, then run the `/TADDM_media/TADDM/setupAix.bin` command on the X Windows environment.
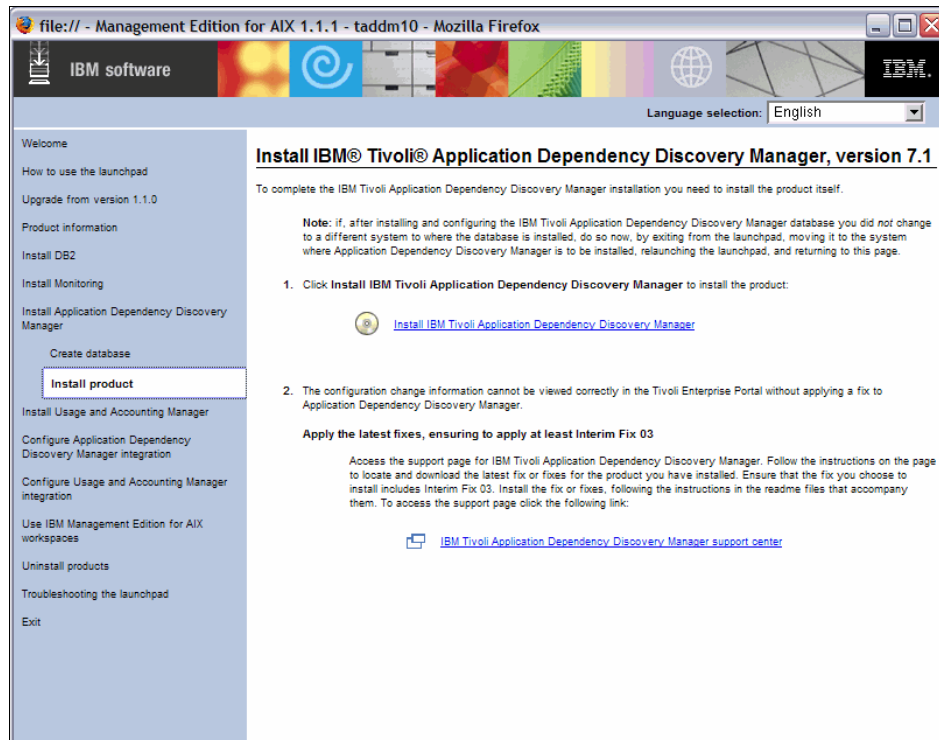
Figure 3-52 Launchpad to install TADDM

6. The next window (Figure 3-53) explains, "`IBM Tivoli Application Dependency Discovery Manager must be installed on a different computer than where its database is installed.`"

   Click **OK** to continue, because we intended to install the database and TADDM on the same server to simplify our configuration.

Figure 3-53   Prompt for database and TADDM on a separate server

7. In the next window, provide a directory path for the TADDM setup and click **OK** as shown in Figure 3-54.

Figure 3-54   Selecting TADDM setup directory

The next window displays the InstallShield Wizard, which is shown in Figure 3-55.

8. Click **Next**.

Figure 3-55   InstallShield Wizard Welcome panel

The next window shows the license agreement for TADDM. Read the licensing terms, and if you agree to the licensing terms, click **I accept both the IBM and the non-IBM terms** as shown in Figure 3-56.

9. Click **Next**.
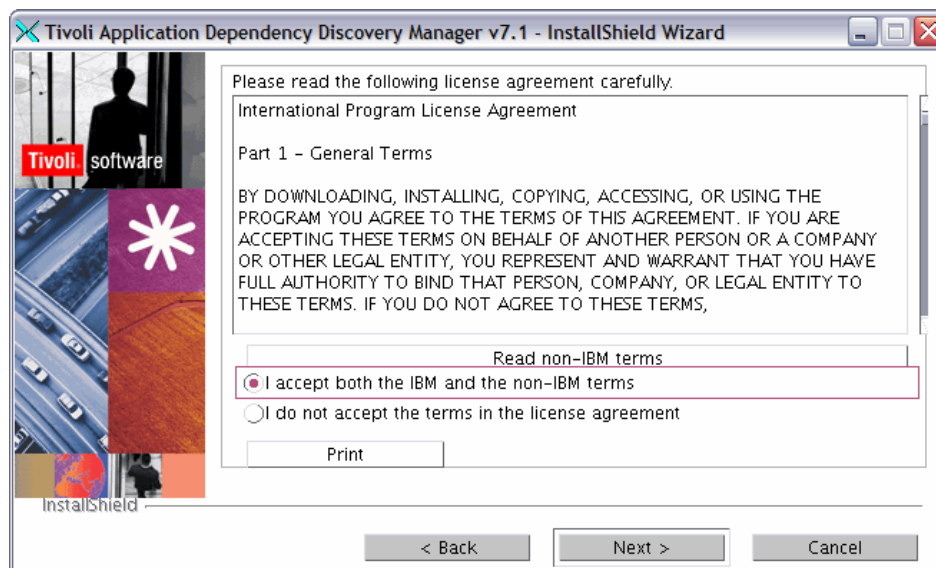
Figure 3-56   License Agreement

In the next window, you will be prompted to enter the name of the directory where you want to install TADDM. We chose `/opt/IBM/cmdb` as shown in Figure 3-57. Click **Next.**
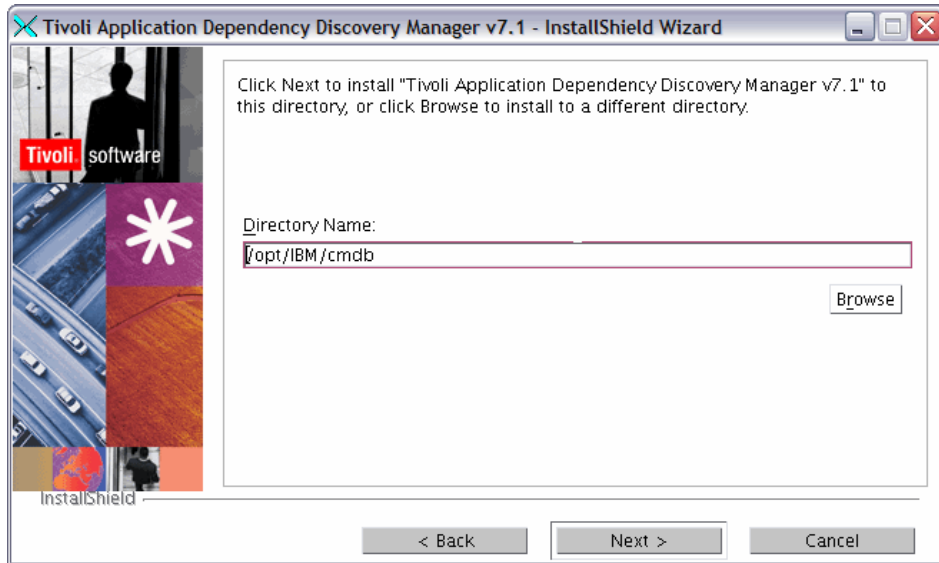
Figure 3-57   Selecting the installation directory for TADDM

In the next window, you can enter the user that you want to start the TADDM Server process. This user must be a non-root user. If this user does not exist, you might request that the InstallShield Wizard create the user account for you by selecting **Create user ID if it does not exist** as shown in Figure 3-58.

10.Click **Next**.
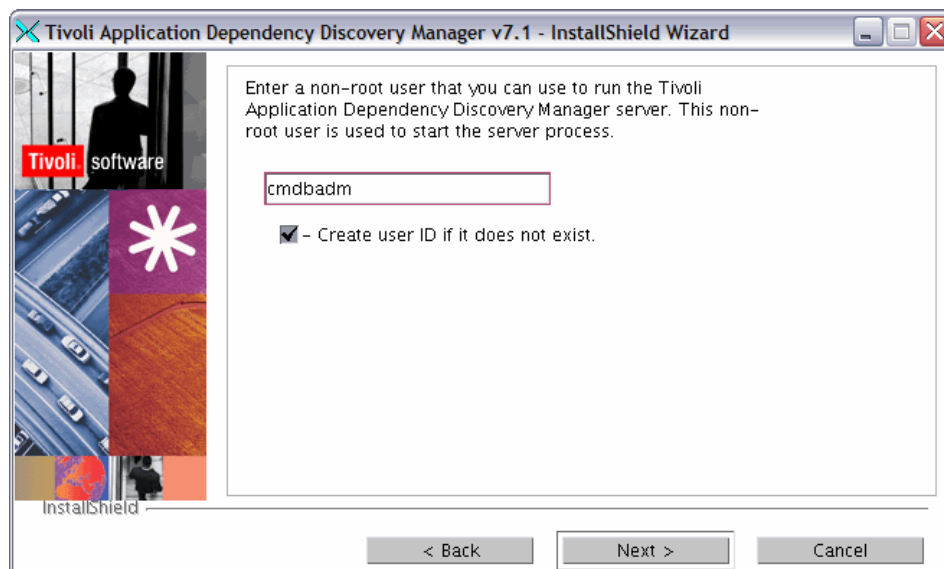
Figure 3-58   Defining a TADDM user

The next window shows installation type options. Choose the installation type: simple or advanced. A simple installation uses default values for a local DB2 database. We chose **Advanced** because we have a database up and running for TADDM as shown in Figure 3-59.

11. Click **Next**.
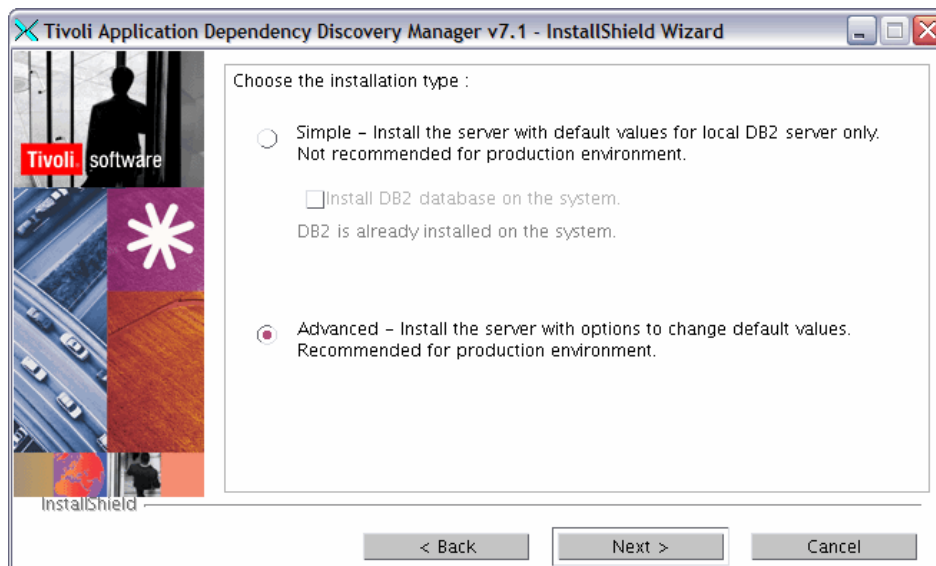
Figure 3-59   Choose the installation type

In the next window, you can select the server type. Select the server type for the TADDM Server that you are installing. We chose **Domain Manager (CMDB) Server**, because we plan to have only one TADDM server in our environment as shown in Figure 3-60.

12. Click **Next**.
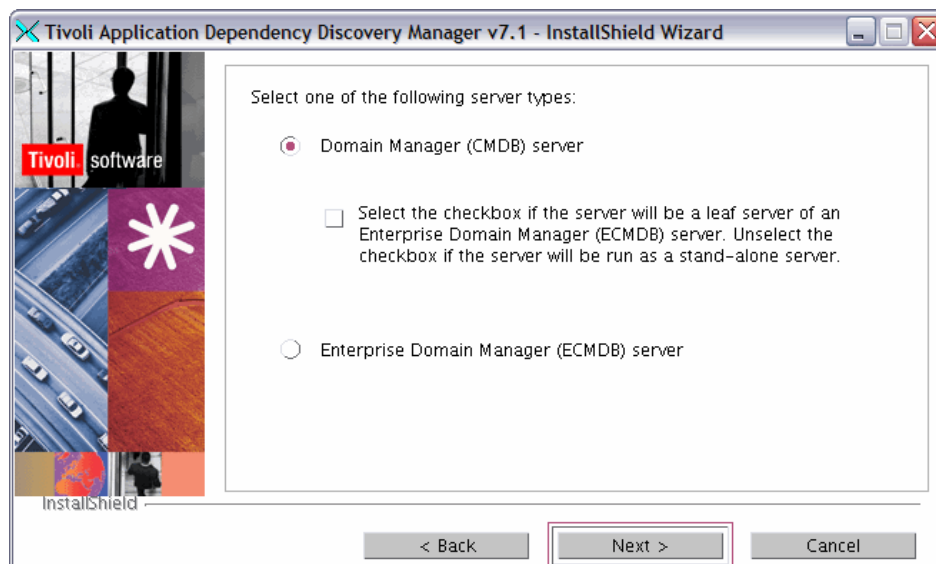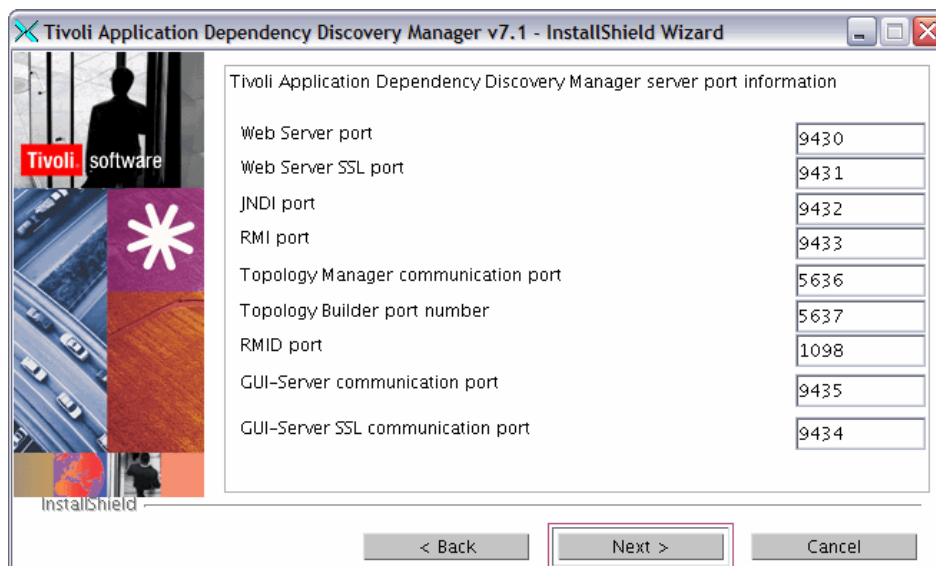
Figure 3-60   Selecting server type

The next window shows the TADDM server port information. Review the default port information and change any port numbers that need to be changed. We accepted all of the default port values as shown in Figure 3-61.

13. Click **Next**.

Figure 3-61   Port information for TADDM Server

In the next window, you can provide the port information for the domain manager server. We took the default values as shown in Figure 3-62.

14.Click **Next**.

Figure 3-62   Port information for domain manager server

In the next window, you can enter the host name for the Remote Method Invocation (RMI) server. We took the default for the host name for the RMI server. Next, select the platform binaries that you want installed. Choose **AIX,** because we are not installing the original TADDM, but the TADDM that is part of Management Edition (ME) for AIX.

15.To start this TADDM Server when the system is started, select **Start the server at system boot**.

16. To start the server after the installation of TADDM is complete, select **Start the server after install** as shown in Figure 3-63. Click **Next**.
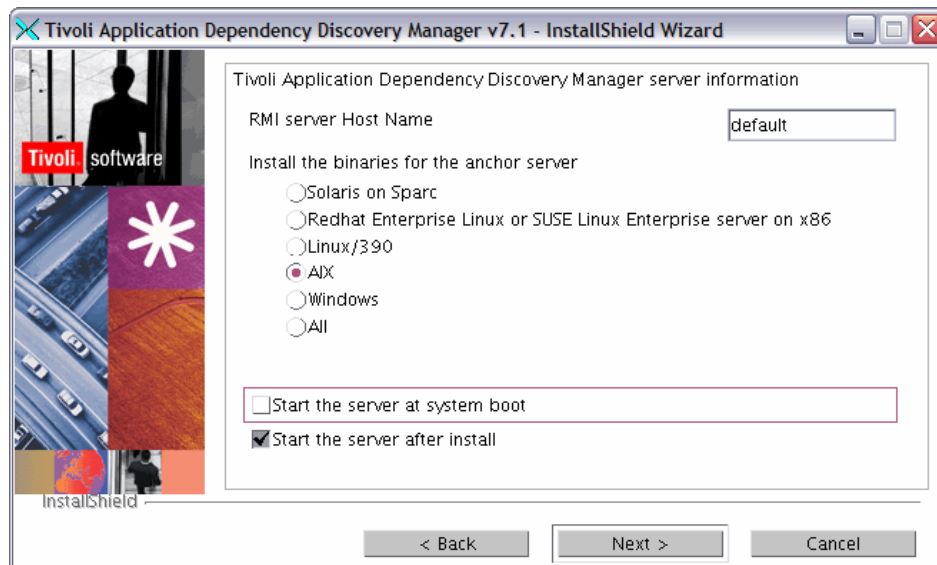
Figure 3-63   Specifying RMI information

The next window, which is shown in Figure 3-64, asks for the host name and port number of the Change and Configuration Management Database (CCMDB). We left this panel blank, because our implementation did not include a CCMDB.

17.Click **Next**.
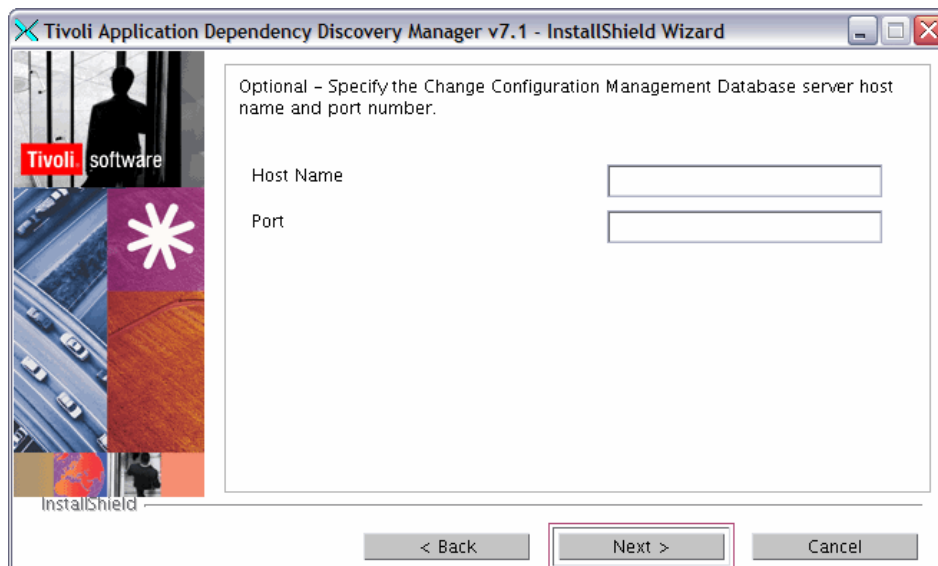
Figure 3-64   Optional CCMDB host name and port

In the next window, which is shown in Figure 3-65, you can select the database type for the TADDM database that you are installing.

18. We are using DB2, so we selected **DB2**.

We chose not to set up the WebSphere Federation Server; therefore, we did not select the Setup WebSphere Federation Server check box.

19. Click **Next**.
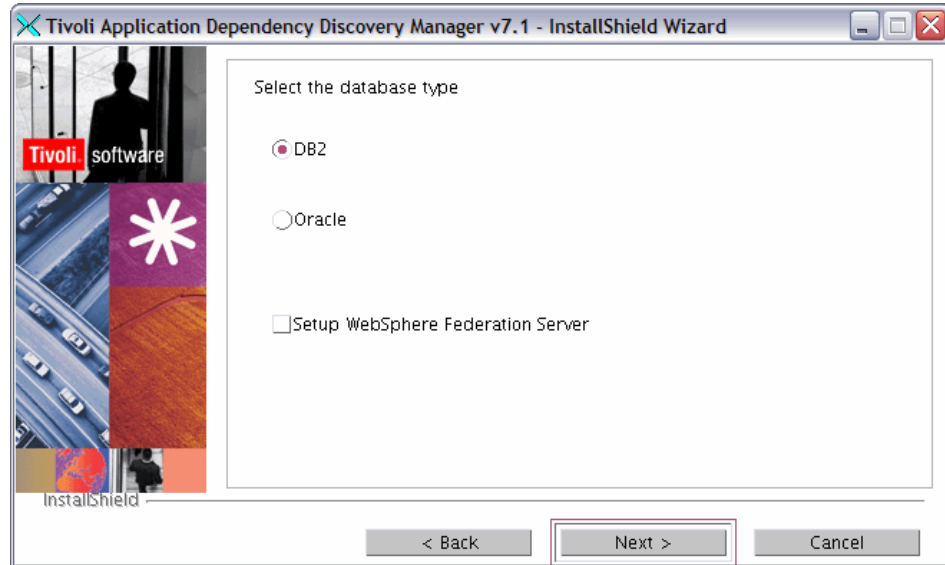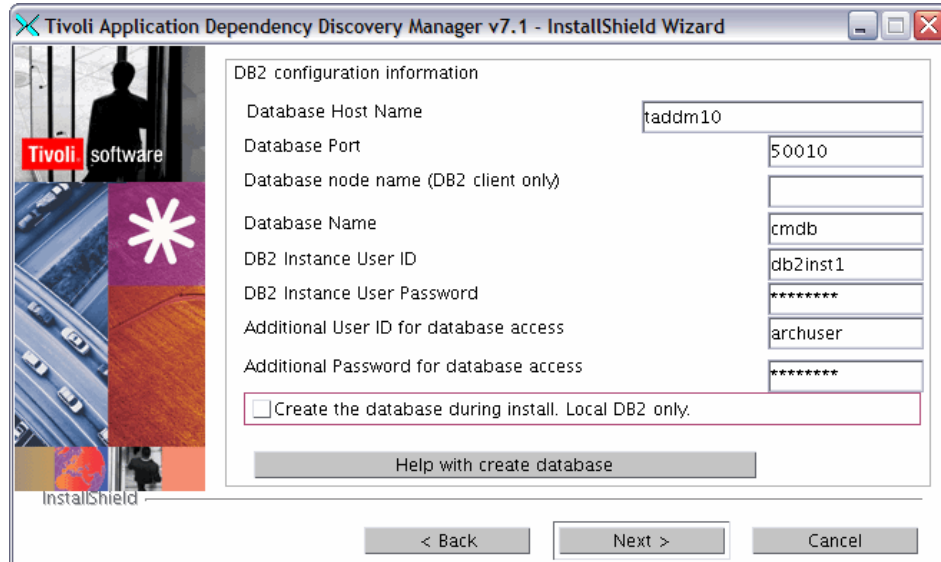
Figure 3-65   Select database type

The next window asks for information about the database. We created our database prior to installing TADDM. Enter the database information here as shown in Figure 3-66.

20. Click **Next**.

Figure 3-66   Database configuration information

In the next window, you can select the option for the user registry that will be used with TADDM.

21. We selected **File Based User Registry** in this configuration as shown in Figure 3-67.

22. Click **Next**.
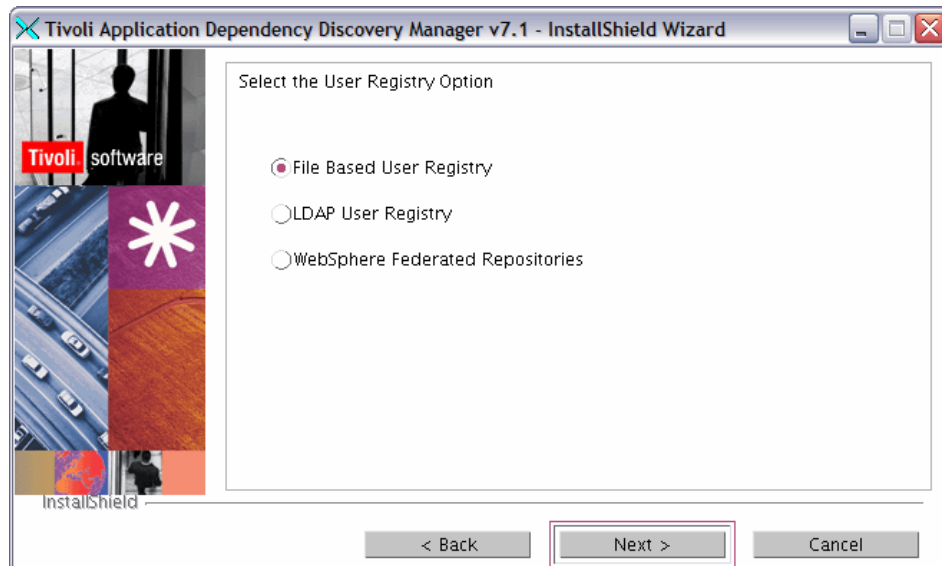
Figure 3-67   Select user registry

The next window shows the summary information.

23. If the information is correct, click **Install** to begin the installation as shown in Figure 3-68.
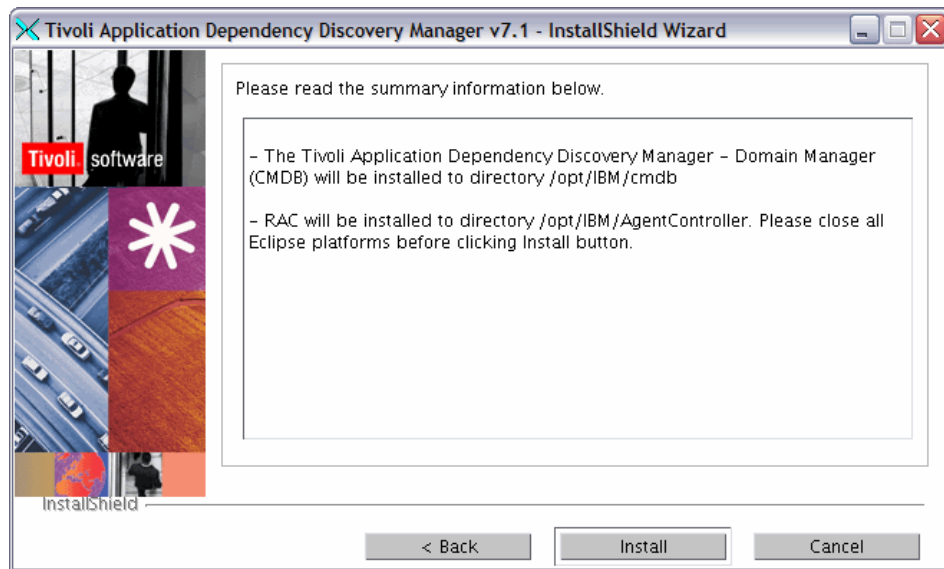
Figure 3-68   Summary information

24. When the installation completes, review the installation summary information, and click **Finish** as shown in Figure 3-69.
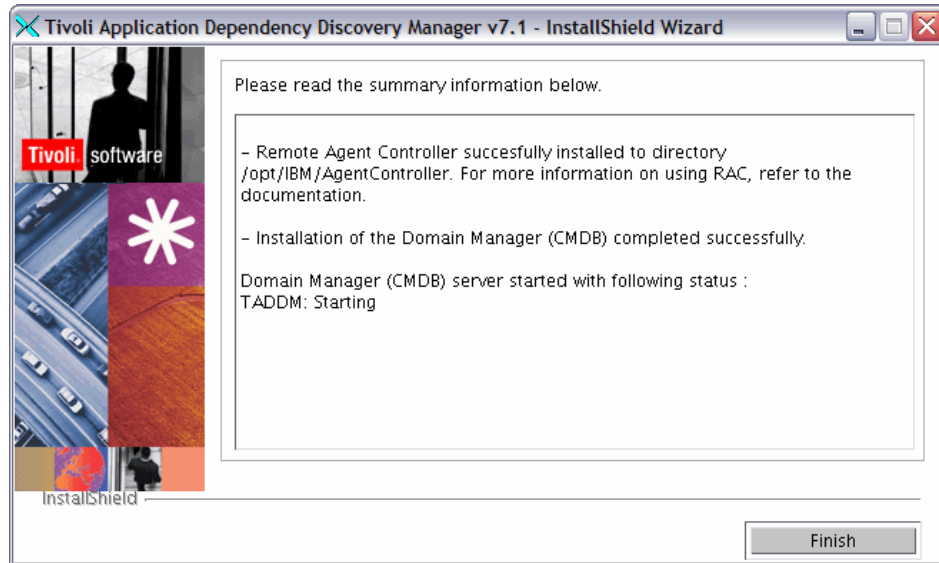
Figure 3-69   Completion of TADDM installation

3.4.5 DNSLookup configuration

If you do not use Domain Name System (DNS) in your environment, edit the `$TADDM_installed_directory/dist/etc/collation.properties` file to disable DNSLookup (In our example, the $TADDM_installed_directory is `/opt/IBM/cmdb.`):

► `com.collation.platform.os.disableDNSLookups=true` (the default is `false`)

► `com.collation.platform.os.disableRemoteHostDNSLookups=true` (the default is `false`)

3.5 Installing and configuring IBM Tivoli Usage and Accounting Manager

In this section, we provide the steps to install Tivoli Usage and Accounting Manager:

3.5.1 Our lab environment

Our lab environment for this installation consists of the Tivoli Usage and Accounting Manager server, three AIX Client LPARs, and a Virtual I/O Server. The TUAM server has a DB2 database that is installed locally on the system. The client LPARs and the Virtual I/O Server are configured with AIX Advanced Accounting, and then, the data collector is installed on each of the systems as shown in Figure 3-70.
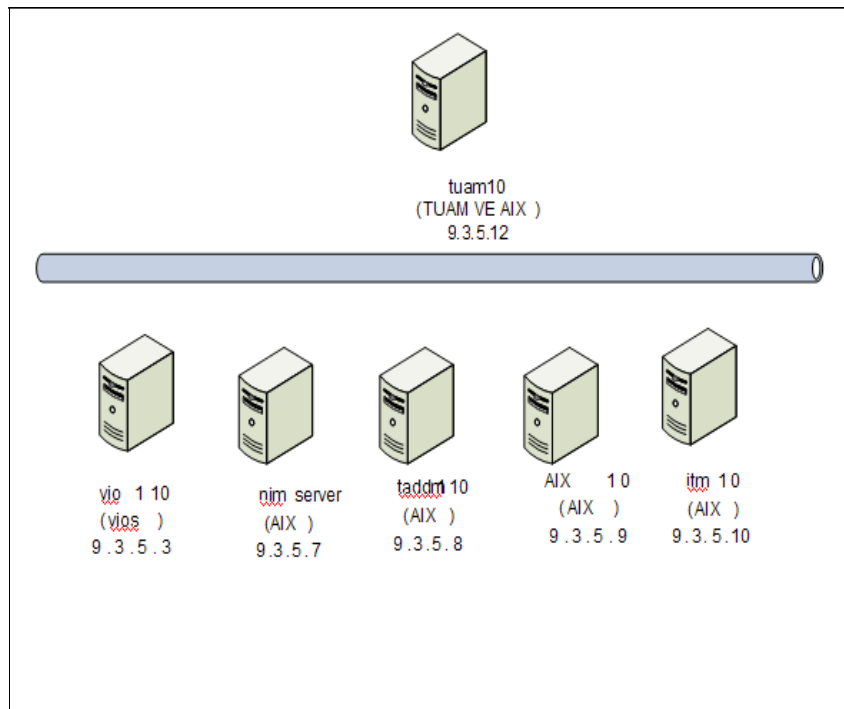

Figure 3-70   TUAM lab set up

3.5.2 Installation process overview

At a high level, the installation process for IBM Tivoli Usage and Accounting Manager Virtualization Edition consists of these tasks:

► Install the application server on the system.

► Create the database, and perform the initial configuration and database initialization.

► Install the Business Intelligence Reporting Tools (BIRT) reporting server. We discuss this topic in 4.4.4, "Installing Business Intelligence and Reporting Tools (BIRT)" on page 250.

► Install data collection in the target systems.

3.5.3 Installing the Application server

To complete the installation of IBM Tivoli Usage and Accounting Manager:

1. Log in to the server as root.

2. Change your directory to the location that contains the Launchpad program. In our example, we have mounted and copied all of the installation files that are needed to a NFS server and we are mounting remotely over `/mnt`. Run `launchpad.sh`, and the Launchpad program appears.

3. The installation of DB2 Enterprise Server Edition is a prerequisite to installing IBM Tivoli Usage and Accounting Manager. We provide the instructions for installing DB2 in 3.2, "Installing and configuring DB2" on page 34. If you have not already installed DB2, you can install it at this time either by following the steps provided in that section, or simply by choosing to install DB2 at the same time that you set up the database in the Launchpad. In the Launchpad, select **Install Usage and Accounting Manager** → **Create Database** as shown in Figure 3-71.
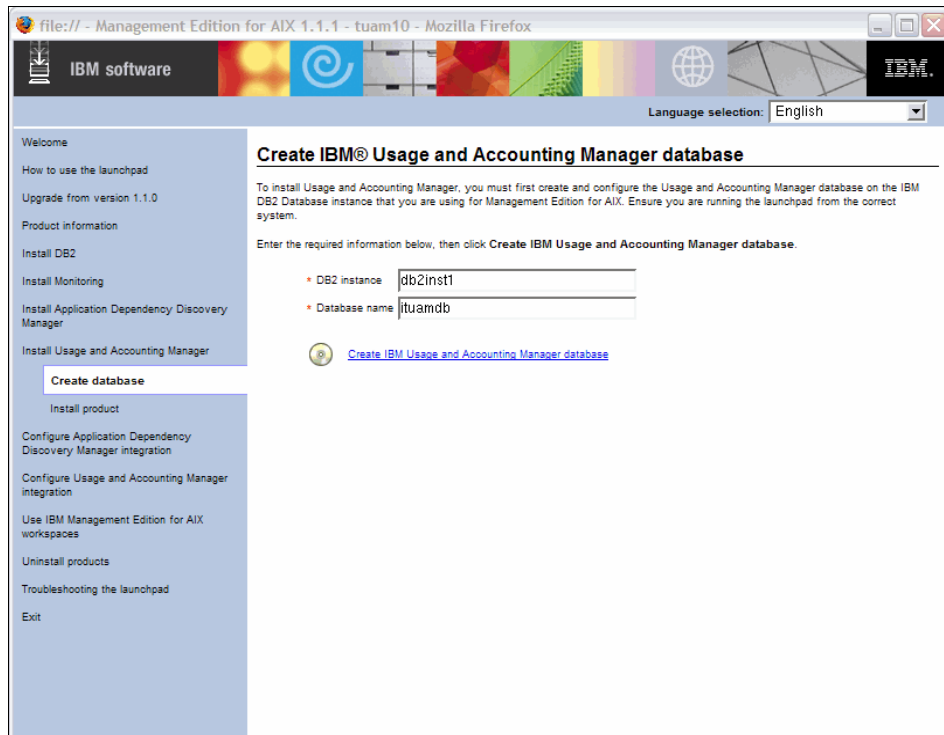
Figure 3-71   Create database

The database is created, and the log appears as shown in Figure 3-72. Click **Close** to continue.
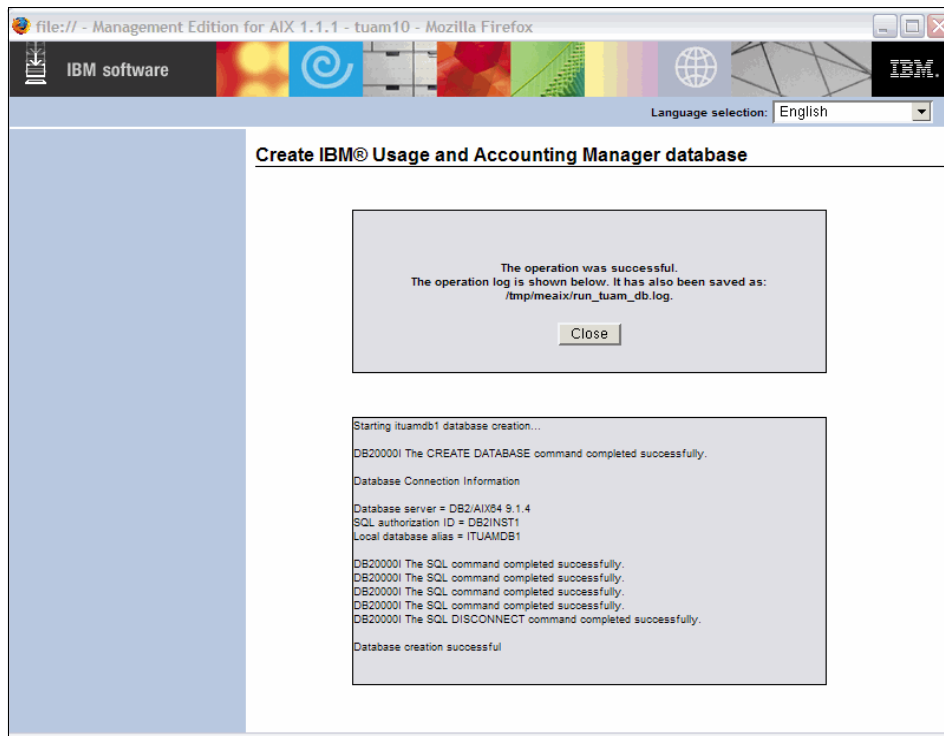
Figure 3-72   Successful database completion message

4. In the Launchpad, select **Install Usage and Accounting Manager** → **Install Product** → **Install IBM Usage and Accounting Manager** as shown in Figure 3-73 on page 122.
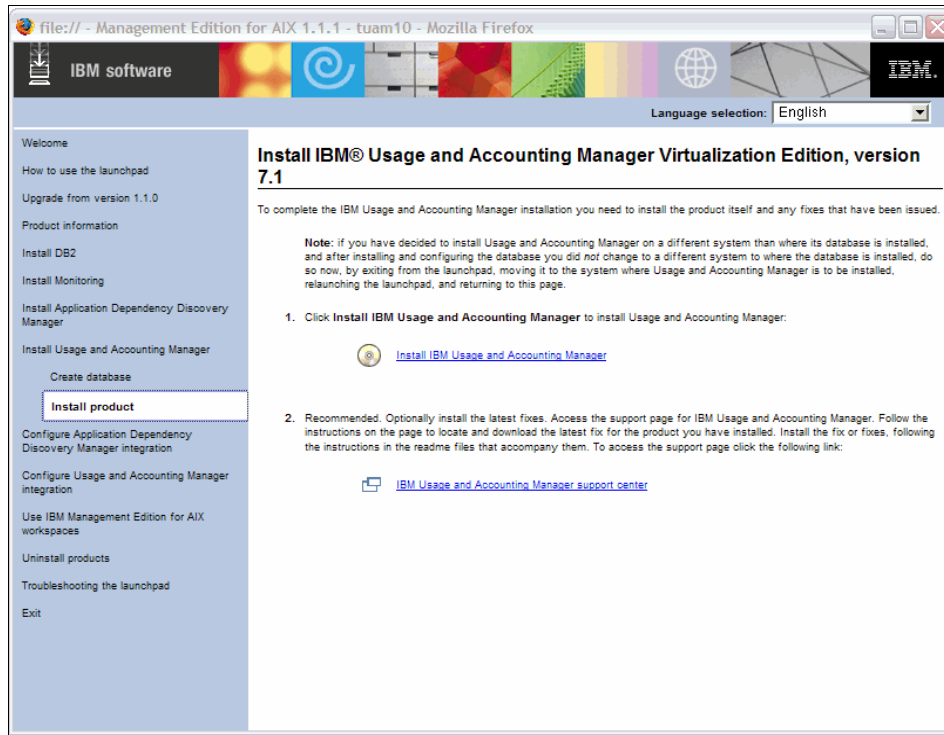
Figure 3-73   IBM Tivoli Usage and Accounting Manager Launchpad Installation

5. After choosing to install IBM Tivoli Usage and Accounting Manager, you will be prompted for the installation path as shown in Figure 3-74. Enter your path and click **OK.**
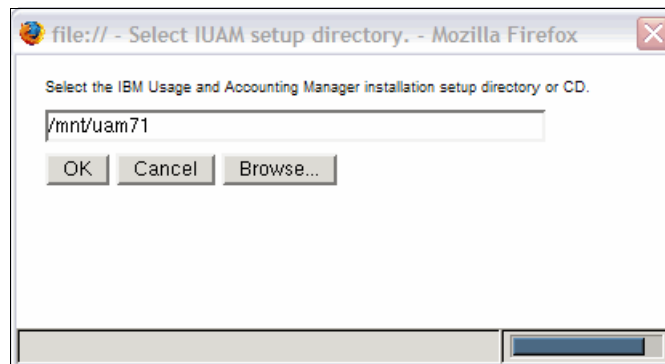
Figure 3-74   CD prompt

6. The InstallSheild Wizard for Usage and Accounting Manager Virtualization Edition starts as shown in Figure 3-75. Choose **Next** to continue.
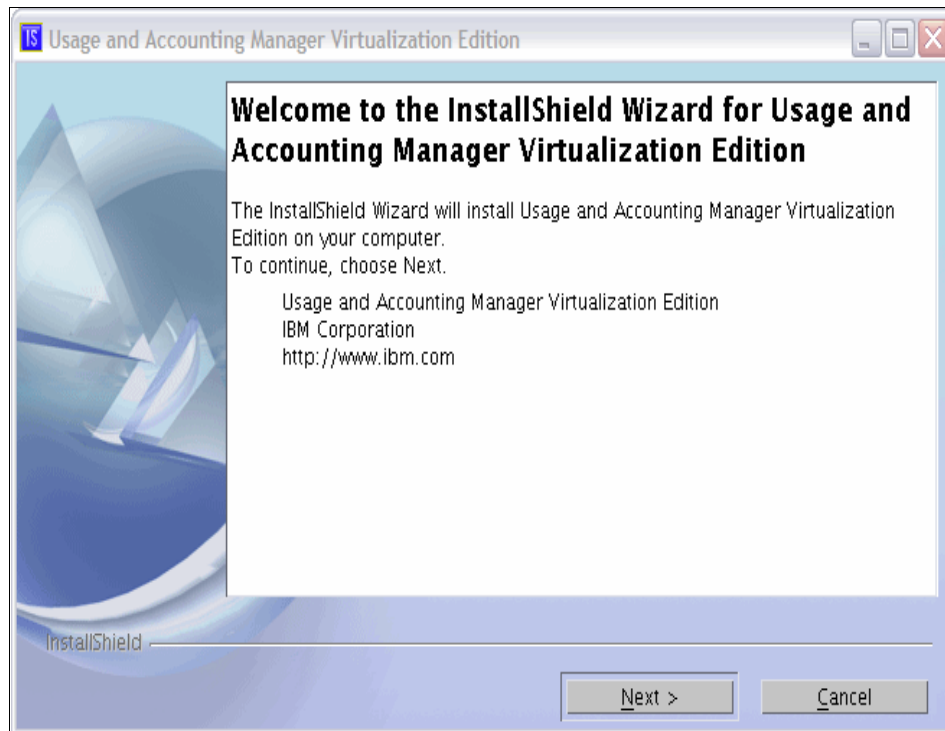

Figure 3-75 InstallShield Wizard for Usage and Accounting Manager

You will be presented with the license terms and agreement as displayed in Figure 3-76.

7. Read the terms and agreement, and if you accept these terms, choose **I accept both the IBM and the non-IBM terms** and then click **Next** to continue.
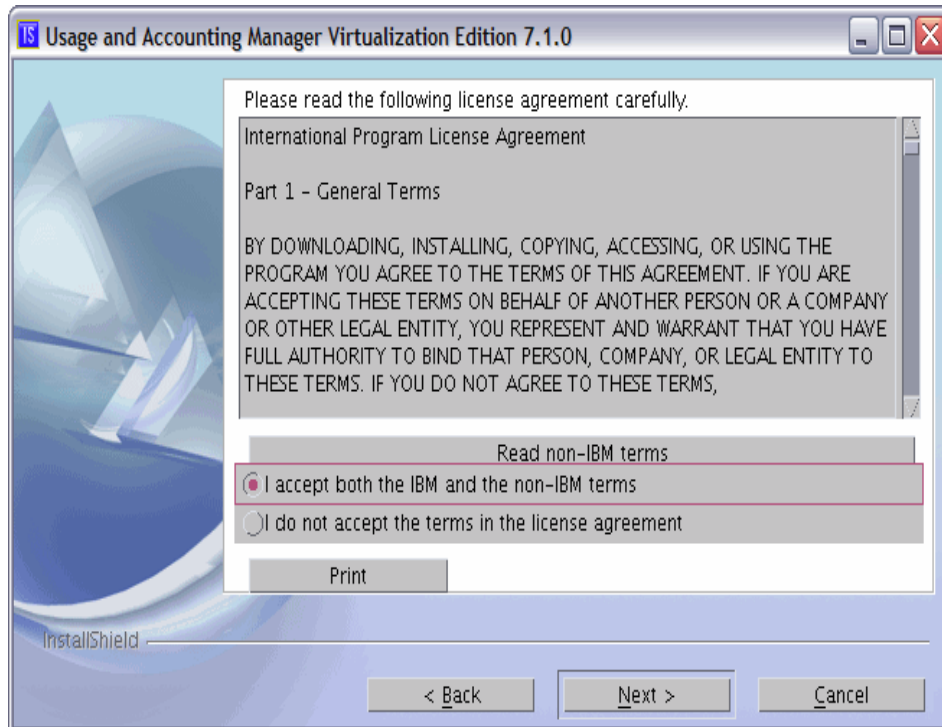

Figure 3-76   IBM Tivoli Usage and Accounting Manager Virtualization Edition License Agreement

You will then see the InstallShield summary as displayed in Figure 3-77. The InstallShield summary gives details of exactly what will be installed on your system and the location and size of the installation.

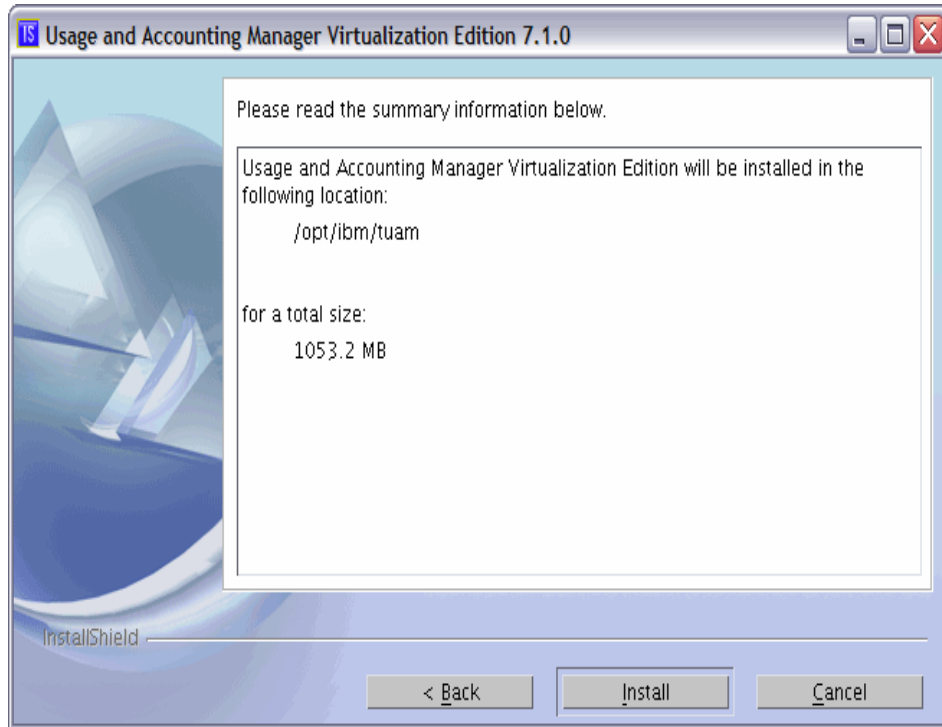8. Click **Install** to continue or **Cancel** to quit the installation.

Figure 3-77   IBM Tivoli Usage and Accounting Manager Virtualization Edition summary information

While IBM Tivoli Usage and Accounting Manager installs, the InstallShield changes status to update the progress of your installation. The installation takes several minutes and might appear stalled on a couple of the windows for several minutes prior to continuing.

Figure 3-78 is an example of one of the installation status updates. You can cancel the installation at any point during the install by selecting **Cancel** while the InstallShield is running.
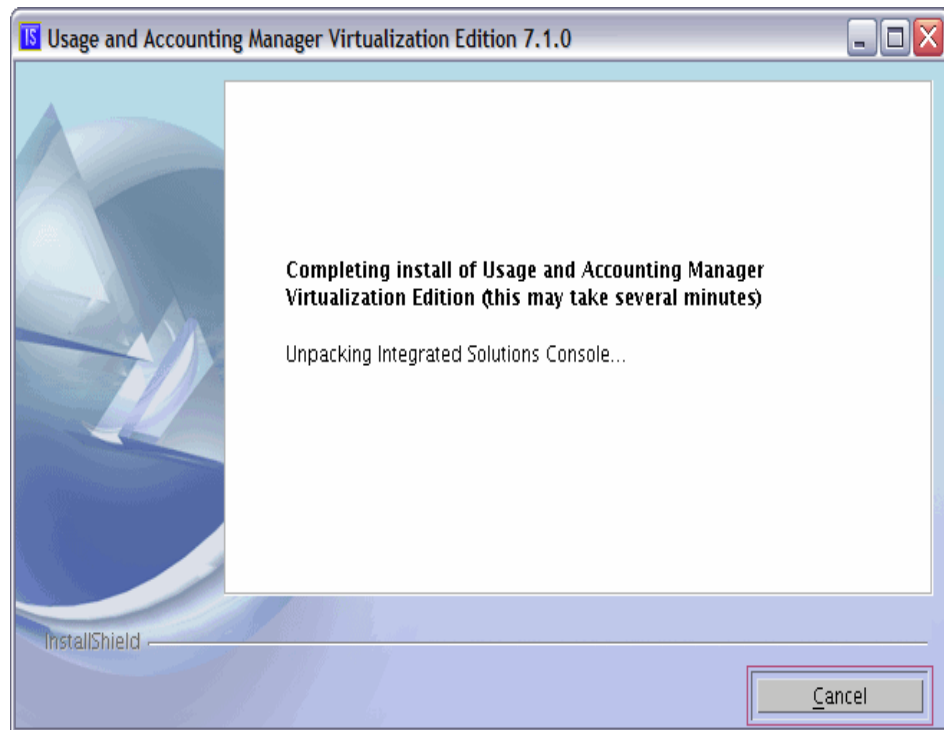

Figure 3-78   IBM Tivoli Usage and Accounting Manager Virtualization Edition InstallShield status

After the installation completes, the InstallShield completion panel appears as shown in Figure 3-79. If there are any issues with the installation, you can obtain the install log at:

```
/opt/ibm/tivoli/common/AUC/logs/install
```

9. If the installation succeeds, select **Finish** to exit the wizard.

Figure 3-79   IBM Tivoli Usage and Accounting Manager Virtualization Edition installation completion message

Verifying the application server

After the application server is installed, run the Tivoli Usage and Accounting Manager console application:

1. Change directory to the embedded WebSphere Application Server directory:

```
/opt/ibm/tuam/ewas/profiles/AppSrv01/bin
```

2. In this directory, run the script to first stop the embedded WebSphere Application Server and then to restart the server using the **stopServer.sh** and **startServer.sh** commands. The application name for TUAM is `server1`. Example 3-16 on page 128 shows an example of successfully running the commands.

Example 3-16   TUAM embedded WAS server restart

```
tuam10:/opt/ibm/tuam/ewas/profiles/AppSrv01/bin ] ./stopServer.sh server1
ADMU0116I: Tool information is being logged in file
           /opt/ibm/tuam/ewas/profiles/AppSrv01/logs/server1/stopServer.log
ADMU0128I: Starting tool with the AppSrv01 profile
ADMU3100I: Reading configuration for server: server1
ADMU3201I: Server stop request issued. Waiting for stop status.
ADMU4000I: Server server1 stop completed.
tuam10:/opt/ibm/tuam/ewas/profiles/AppSrv01/bin ] ./startServer.sh server1
ADMU0116I: Tool information is being logged in file
           /opt/ibm/tuam/ewas/profiles/AppSrv01/logs/server1/startServer.log
ADMU0128I: Starting tool with the AppSrv01 profile
ADMU3100I: Reading configuration for server: server1
ADMU3200I: Server launched. Waiting for initialization status.
ADMU3000I: Server server1 open for e-business; process id is 692414
```

3. Open a Web browser and direct it to the Integrated Solution Console URL:

   ```
   http://tuam10:11052/ibm/console
   ```

4. After you log in, the Welcome page appears (Figure 3-80). Verify that the product information matches the version that you have installed.
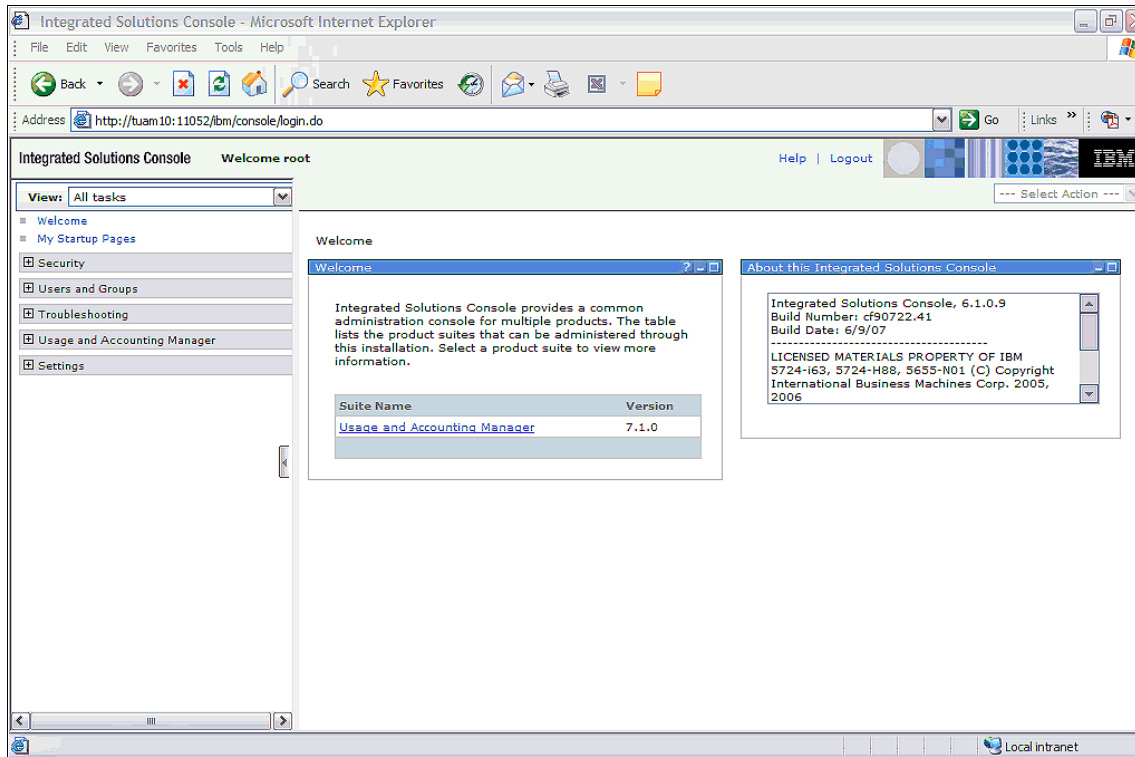

Figure 3-80   Integrated Solutions Console Welcome page

3.5.4 Database configuration

After creating the initial database and installing the application server, the next step is to configure the database server to connect correctly to the application server. Tivoli Usage and Accounting Manager requires a database with a page size of at least 16 KB. This size is not the default. In order to change it:

1. Log on to the system containing the DB2 UDB database as the database user, `db2inst1`. Query the DB2 system tables to determine whether a large enough buffer pool exists. The SQL statement and results are in Example 3-17 on page 130.

Example 3-17   Connecting to database and listing buffer pool

```
$ tuam10:/mnt/launchpad ] su - db2inst1
$ db2 connect to ituamdb

   Database Connection Information

 Database server        = DB2/AIX64 9.1.4
 SQL authorization ID   = DB2INST1
 Local database alias   = ITUAMDB

$ db2 "SELECT PAGESIZE, SUBSTR(BPNAME,1,20) AS BUFFERPOOL FROM
SYSCAT.BUFFERPOOLS"

PAGESIZE    BUFFERPOOL
----------- --------------------
       4096 IBMDEFAULTBP
```

2. Define a buffer pool with a page size larger than the 4 KB default by using the CREATE BUFFERPOOL command. The SQL syntax is shown in Example 3-18.

Example 3-18   Creation of larger buffer pool

```
$ db2 "create bufferpool BP32K size 1000 pagesize 32k"
DB20000I  The SQL command completed successfully.
```

3. Create a regular table space that uses the larger buffer pool as shown in Example 3-19.

Example 3-19   Creation of a regular table space

```
$ db2 "create tablespace USERSPACE2 PAGESIZE 32K managed by system
using ('/home/db2inst1/db2inst1/NODE0000/SQL00002/SQLT004') bufferpool
BP32K"
DB20000I  The SQL command completed successfully.
```

4. Create a temporary table space that uses the larger buffer pool as shown in Example 3-20. Ensure that the temporary parameter is used when creating this table space.

Example 3-20   Creation of temporary table space

```
$ db2 "create temporary tablespace TEMPSPACE2 PAGESIZE 32K managed by
system using ('/home/db2inst1/db2inst1/NODE0000/SQL00002/SQLT005')
bufferpool BP32K"
DB20000I  The SQL command completed successfully.
```

5. Restart the database instance. Use the **LIST APPLICATION** command to check where the are any processes that are still using the database. The restart process is shown in Example 3-21.

Example 3-21   Restarting the database manager

```
$ db2 list application
SQL1611W  No data was returned by Database System Monitor.
$ db2stop
03/04/2009 10:18:51    O   O   SQL1064N  DB2STOP processing was successful.
SQL1064N  DB2STOP processing was successful.
$ db2start
03/04/2009 10:19:03    O   O   SQL1063N  DB2START processing was successful.
SQL1063N  DB2START processing was successful.
tuam10:/opt/ibm/tuam/ewas/profiles/AppSrv01/bin ] ./startServer.sh server1
ADMU0116I: Tool information is being logged in file
/opt/ibm/tuam/ewas/profiles/AppSrv01/logs/server1/startServer.log
ADMU0128I: Starting tool with the AppSrv01 profile
ADMU3100I: Reading configuration for server: server1
ADMU3200I: Server launched. Waiting for initialization status.
ADMU3000I: Server server1 open for e-business; process id is 692414
$ db2 list application
Auth Id Application    Appl.      Application Id                                                  DB       # of
        Name           Handle                                                                    Name     Agents
-------- -------------- ---------- ------------------------------------------------------------- -------- -----

DB2INST1 db2jcc_applica 34         9.12.5.9.43682.090304163507                                   ITUAMDB  1
DB2INST1 db2evmg_DB2DET 33         *LOCAL.DB2.090304163444                                       ITUAMDB  1
DB2INST1 db2taskd       32         *LOCAL.DB2.090304163442                                       ITUAMDB  1
DB2INST1 db2stmm        31         *LOCAL.DB2.090304163441                                       ITUAMDB  1
DB2INST1 db2jcc_applica 30         9.12.5.9.43680.090304163440                                   ITUAMDB  1
```

> **Note:** If db2 LIST APPLICATION returns connections to the database, which it does in the second run in Example 3-21 on page 131, you must run **/opt/ibm/tuam/ewas/profiles/bin/stopServer.sh server1** to terminate the connections. If you are running other applications that are accessing DB2, you must also stop those applications.

Configuring the JDBC driver

Tivoli Usage and Accounting Manager requires two JDBC drivers to connect to the DB2 database. The drivers are included with the DB2 base installation, and you need to load them into IBM Tivoli Usage and Accounting Manager using the Integrated Solutions Console. The default path for the JDBC drivers is `/opt/IBM/db2/V9.1/java`, and the two drivers that we need to use are:

► `db2jcc.jar`

► `db2jcc_license_cu.jar`

To configure the JDBC driver:

1. To add the drivers to TUAM, open the Integrated Solutions Console (ISC) menu and select **Usage and Accounting Manager** → **System Maintenance** → **Configuration**. On the Drivers tab, click **New** to define a driver. Enter the path and filename of the DB2 UDB drivers as shown in Figure 3-81. Click **OK**.
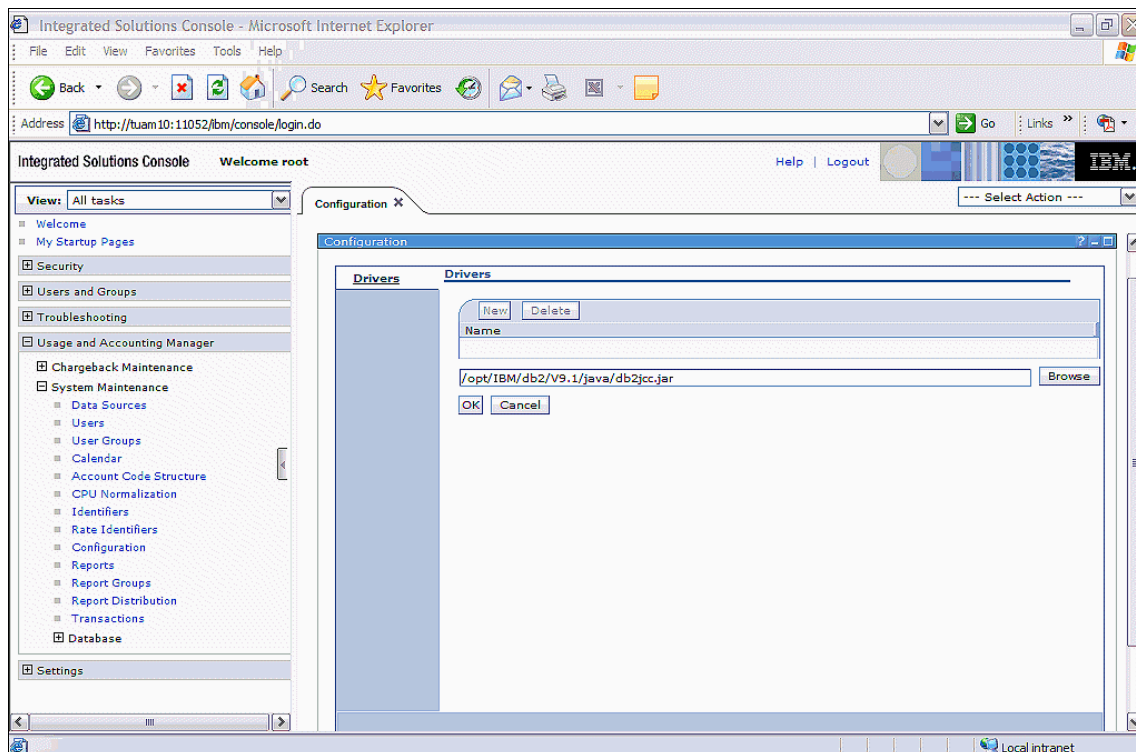
Figure 3-81   Adding JDBC driver db2jcc.jar

2. Repeat the steps to add the second JDBC driver
   `/opt/IBM/db2/V9.1/java/db2jcc_license_cu.jar` to the IBM Tivoli Usage
   and Accounting Manager as shown in Figure 3-82, and click **OK**.
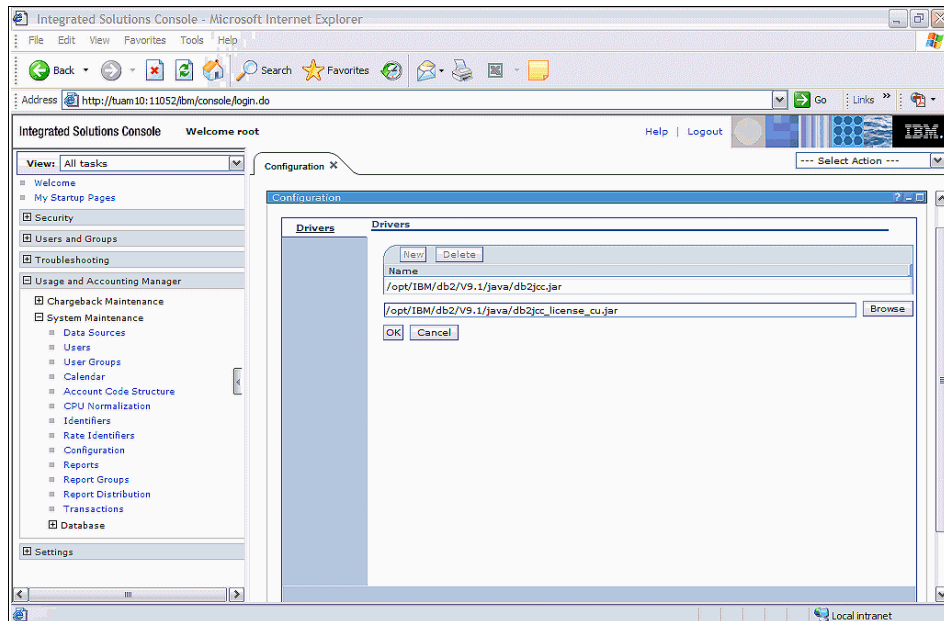
Figure 3-82 Adding JDBC driver db2jcc_license_cu.jar

Configuring the data sources

The second step to accessing the DB2 database from Tivoli Usage and Accounting Manager is to set up the data sources. We must set the database sources for the Server data source at minimum:

1. In the ISC, select **Usage and Accounting Manager** → **System Maintenance** → **Data Sources** on the Server tab, and the database that you created from the Launchpad will be listed. If you do not see your database listed under the Server tab, you can add your database on this same tab by clicking **New** and filling in the fields as shown in Figure 3-83.

2. Then, click **OK** at the bottom of the page.

   If your database is already listed, you need to edit the data source by selecting the view pop-up menu icon next to the data source **EditDataSource**. The same window that is shown in Figure 3-83 appears. Ensure that you change the fields for your environment, including User Name, Password, Host, and Database Type. When complete, click **OK** on the bottom of the page.
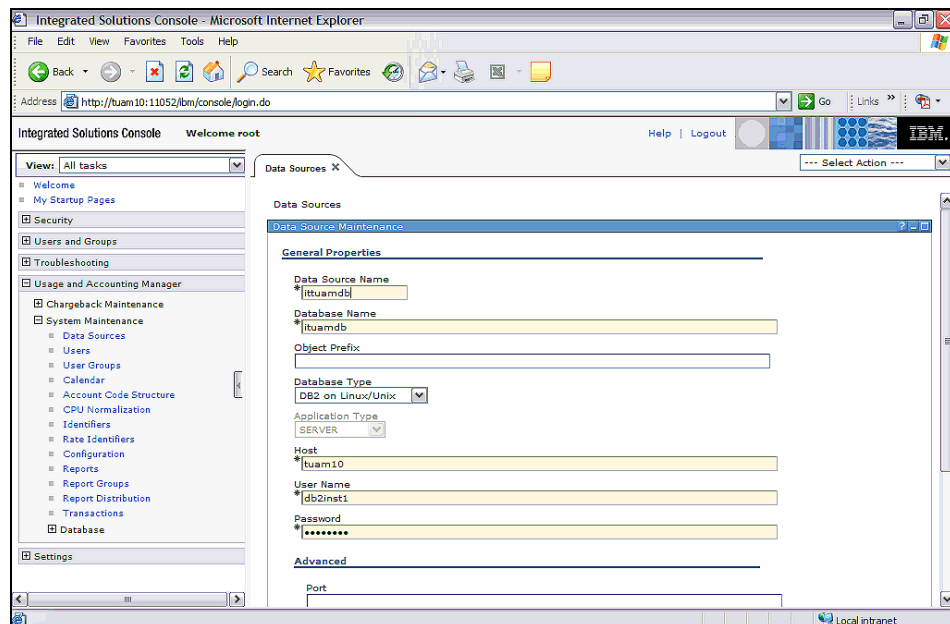
Figure 3-83   Adding or editing a data source

If you use multiple databases to store Usage and Accounting Manager (for example, you have a production database and a development database), you must select the data source for one database as the default for administration, data processing, and Web reporting.

> **Note:** This process is applicable for Server data sources only.

3. In the ISC, click **Usage and Accounting Manager** → **System Maintenance** → **Data Sources** on the Server tab.

   Then, set the Default Admin, Processing, and Reporting data source by selecting the data source name and clicking **Set Admin** from the pop-up menu.

4. Then, select **Set Processing**, and finally, select **Set Reporting** so that all three functions are now set and show on the Server tab with a *yes* in the last three columns as shown in Figure 3-84.
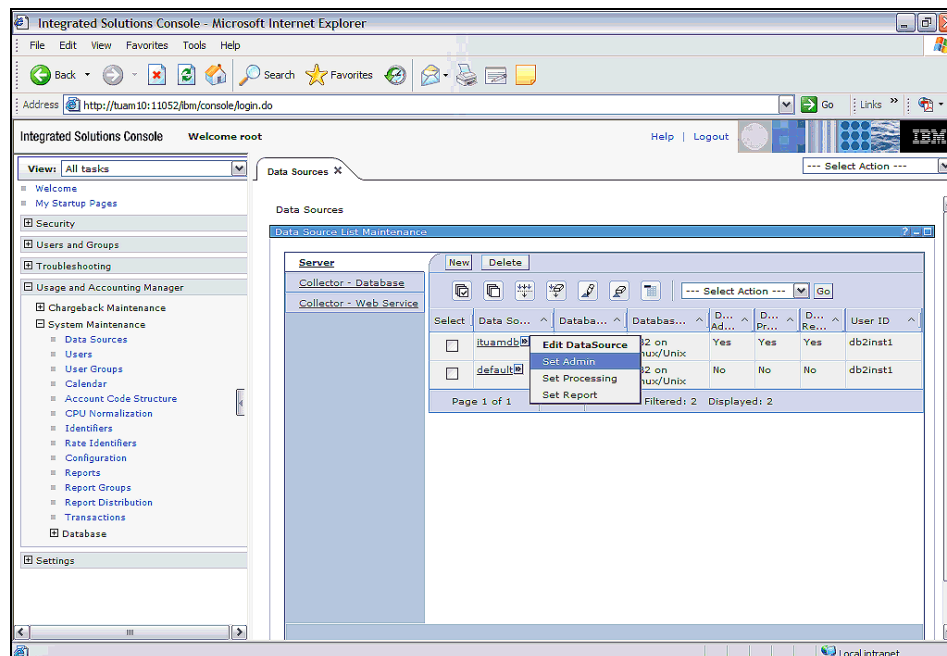
Figure 3-84   Setting the Default Admin, Processing, and Reporting data source

5. After the JDBC drivers have been loaded and the data source properties have all been configured, you must stop and start the embedded WebSphere server for the changes to take affect. As the root user, run:

► `/opt/ibm/tuam/ewas/bin/stopServer.sh server1`
► `/opt/ibm/tuam/ewas/bin/startServer.sh server1`

Verify connectivity to database

We need to test the connection to verify that Tivoli Usage and Accounting Manager is able to communicate to the DB2 database. To test connection in the ISC, go to **Usage and Accounting Manager** → **System Maintenance** → **Data Sources,** click the pop-up menu next to your database name, and select **Edit**. Figure 3-85 appears. Select **Test** at the bottom of the page.
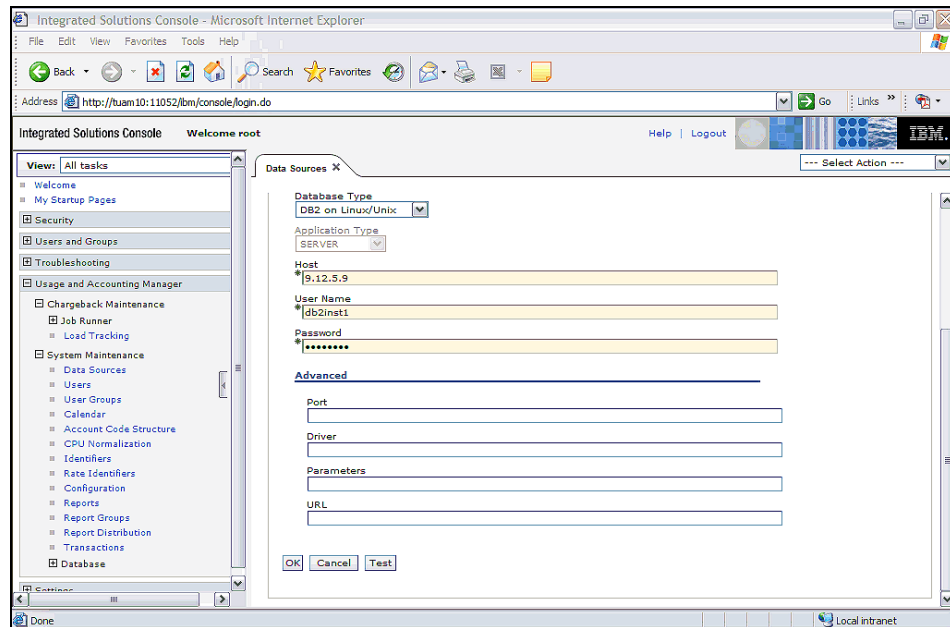
Figure 3-85   Data Source ISC window

You should receive a response that your connection was successful. If you receive any errors, for example, the password is incorrect, correct the issue and try again.

Initialize the database

The steps used to initialize the database are:

1. To initialize the Tivoli Usage and Accounting Manager database using the ISC menu, select **Usage and Accounting Manager** → **System Maintenance** → **Database** → **Initialize Database** as shown in Figure 3-86.

2. Click **Initialize Database** to initialize the database.
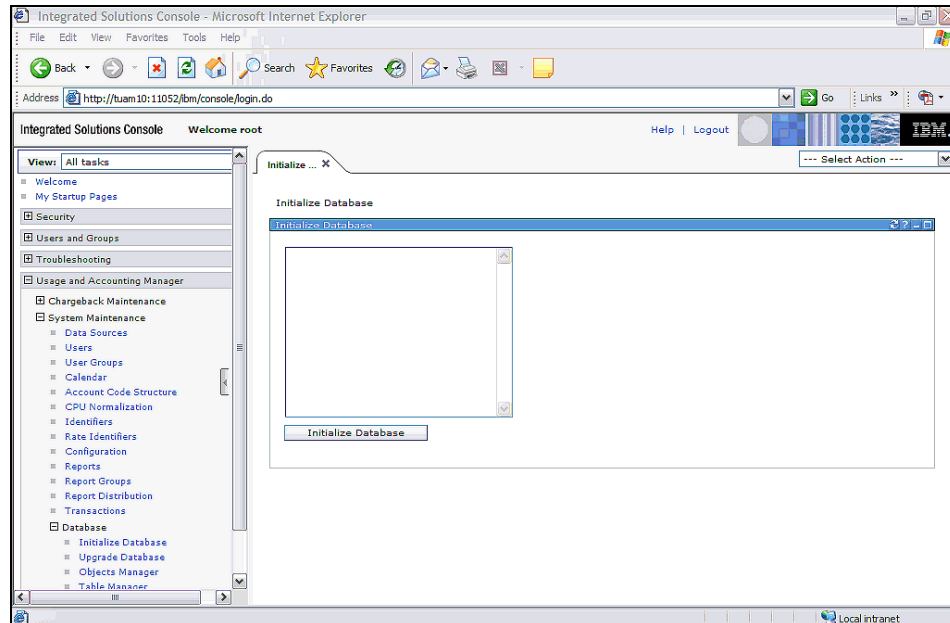
Figure 3-86   Initialize database

The confirmation message that is shown in Figure 3-87 appears. Verify that the database name and server match the system that you want to initialize.

3. Click **Yes** to continue to initialize the database.
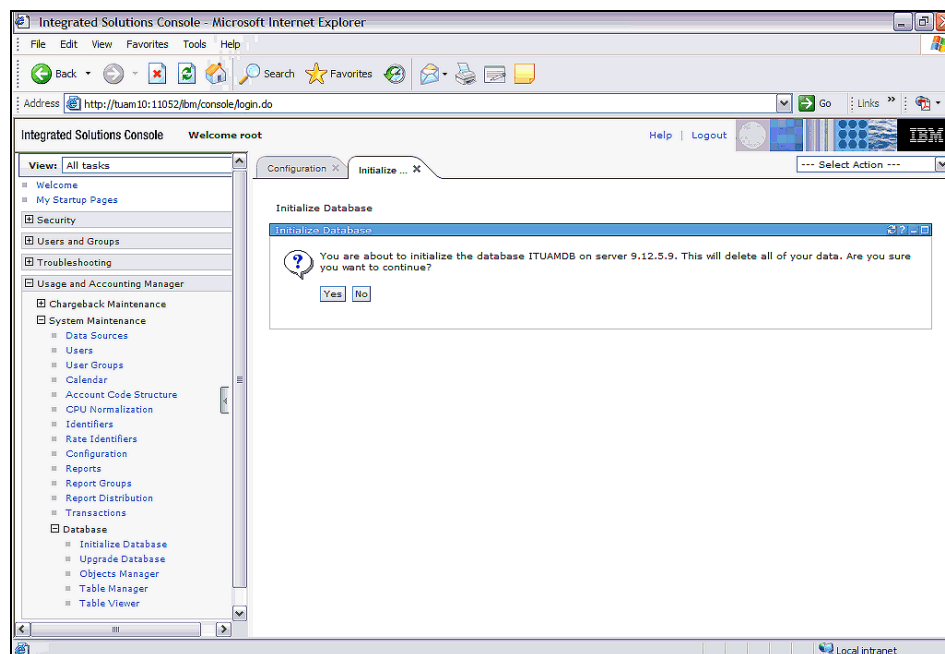
Figure 3-87   Confirmation window for database initialization

The results of the initialization appear.

4. Confirm that the database initialized successfully. If there were any issues, review the log that is present, correct the issue, and try again. A successful initialization log is shown in Figure 3-88.
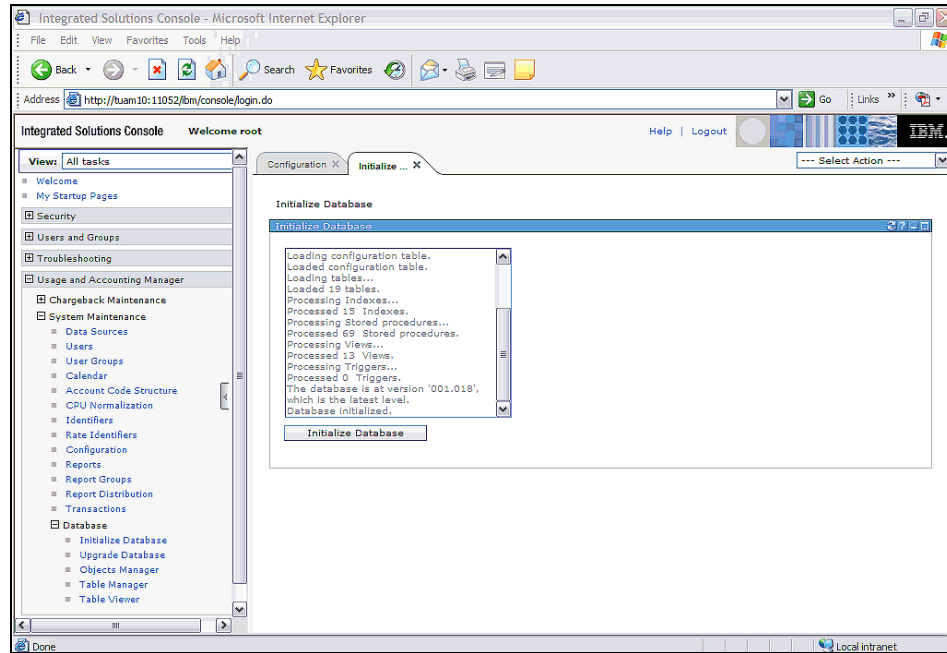
Figure 3-88   Results of database initialization

3.5.5 Set up the processing directories

To set up the processing directories:

1. After you install Usage and Accounting Manager, you must set up directories to store the required files that are listed when you select ISC **Configuration** → **Processing** and **Configuration** → **Reporting**. The Configuration tab Processing page is shown in Figure 3-89 on page 140. In order for a remote system to process or send files to these directories, that system must have read-write permissions on these directories:

   – Job Files Path

     Location of job files or the XML file that defines the process of collecting data and loading it into Usage and Accounting Manager database.

   – Sample Job Files Path

     Directory that contains the sample job files.

– Job Log Files Path

The job logs provide processing results for each step that is defined in the job file. If a warning or failure occurs during processing, the file indicates at which point the warning or failure occurred.

– Collector Log Files Path

Directory that will contain the collector log files.
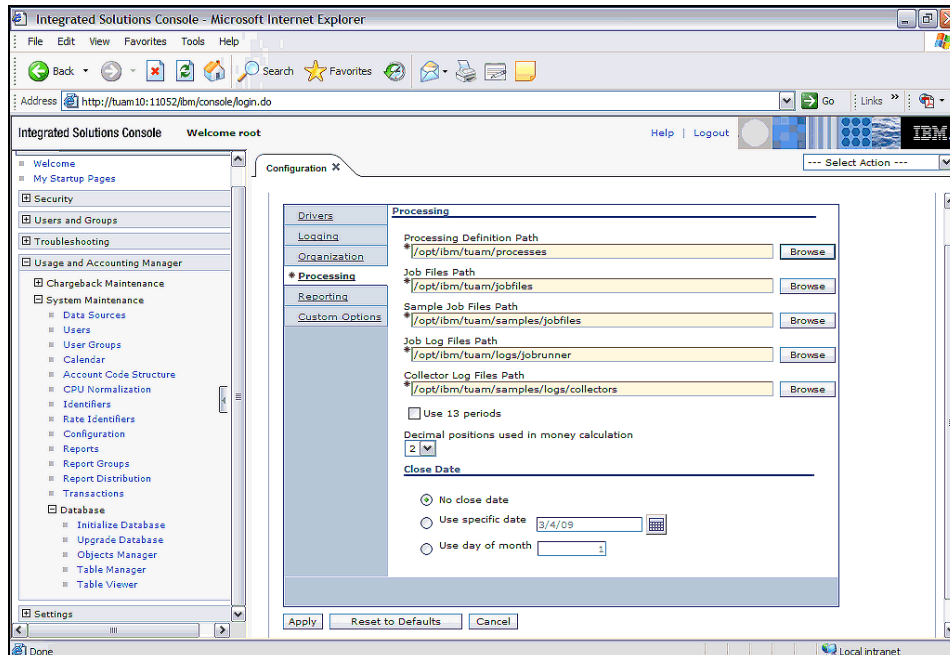
Figure 3-89   ISC Configuration Processing

2. The directories that are on the Configuration tab Reporting page need to be changed from the default directories for Tivoli Usage and Accounting Manager Virtualization Edition. The paths are:

– Path to Standard Reports Folder

Directory that contains Usage and Accounting standard reports. There are five templates included for BIRT with IBM Tivoli Usage and Accounting Manager VE.

– Path to Custom Reports Folder

Directory that contains custom reports. You create custom reports with your report developer either independently or by using one of the standard reports as a template.

– Path to Published Reports Folder

Directory that contains published reports. Published reports are reports that have been saved with the data that was generated at the time that the reports were run.

The type of reporting supported with Tivoli Usage and Accounting Manager Virtualization Edition is BIRT reporting, which is an Open Source reporting tool. We need to set each of the reporting directories to point to the BIRT directories as shown in Figure 3-90.
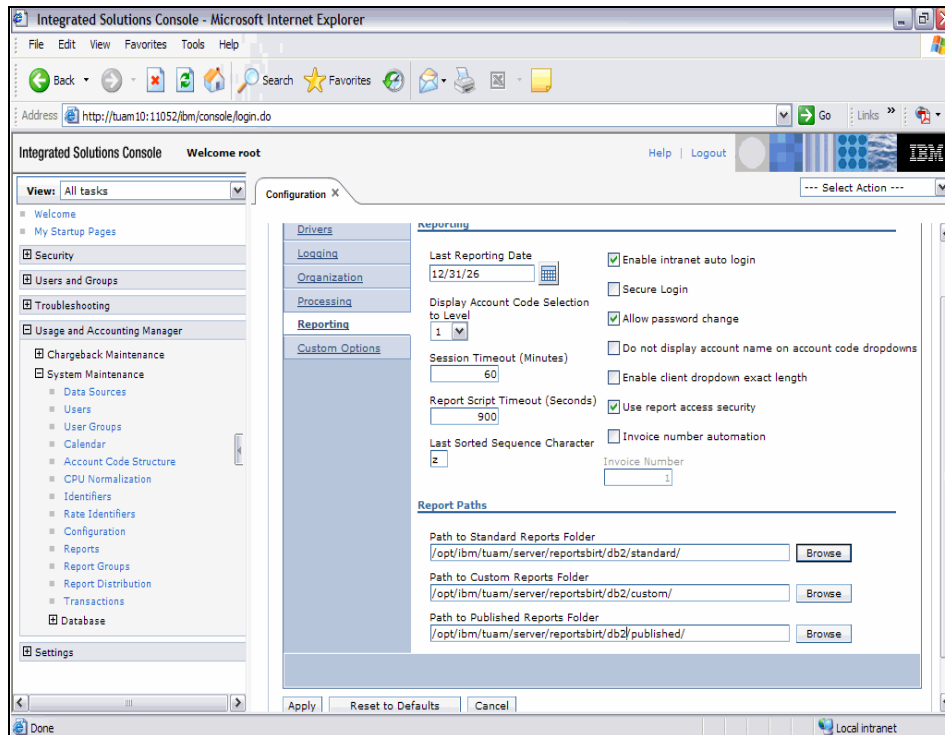
Figure 3-90   ISC Configuration tab Reporting page

3.5.6 Run sample jobs

Sample jobs allow you to validate the installation:

1. Before performing any additional configuration in Tivoli Usage and Accounting Manager, run the provided samples:

   ```
   /opt/ibm/tuam/bin/RunSamples.sh | tee RunSamples.log
   ```

   You can ignore warning messages pertaining to Simple Mail Transfer Protocol (SMTP), because we have not configured any of these jobs to use

SMTP. To check the results, from the ISC menu, select **Usage and Accounting Manager** → **Chargeback Maintenance - Job Runner** → **Log Files** as shown in Figure 3-91.
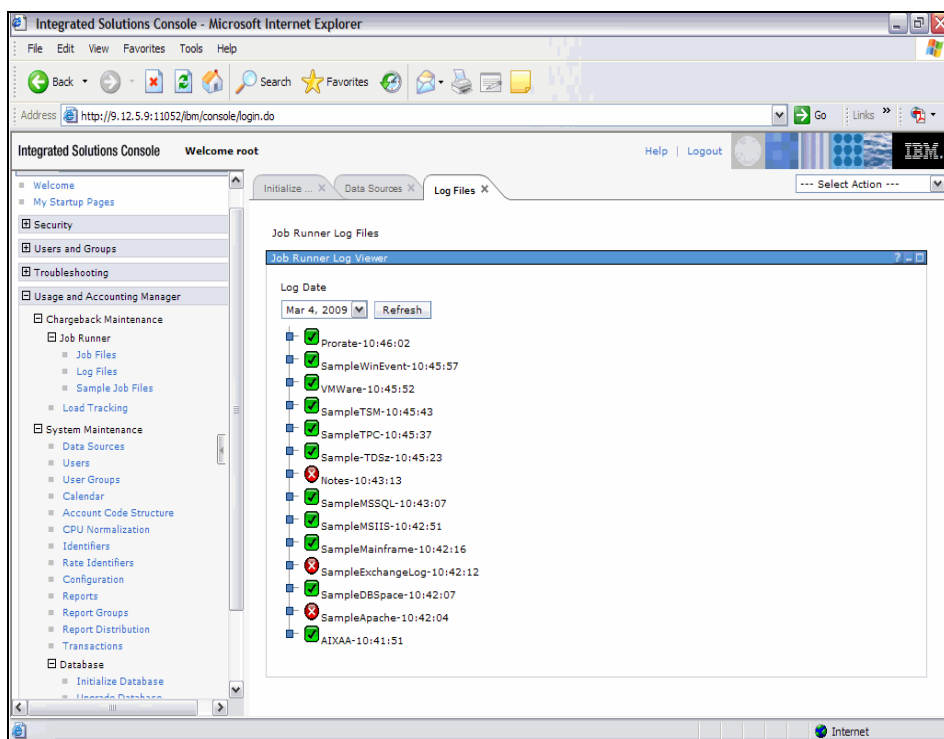
Figure 3-91 Sample log output

Other errors are caused by not installing the enterprise collector pack (ecp). To perform a final database cleanup, re-initialize it as described in the previous section, "Initialize the database" on page 137.

Alternatively, you can unload the data using the ISC menu selections **Usage and Accounting Manager** → **Chargeback Maintenance** → **Load Tracking** as shown in Figure 3-92. Clear the check box for the **End Date** filter and mark all by clicking the Select check boxes.
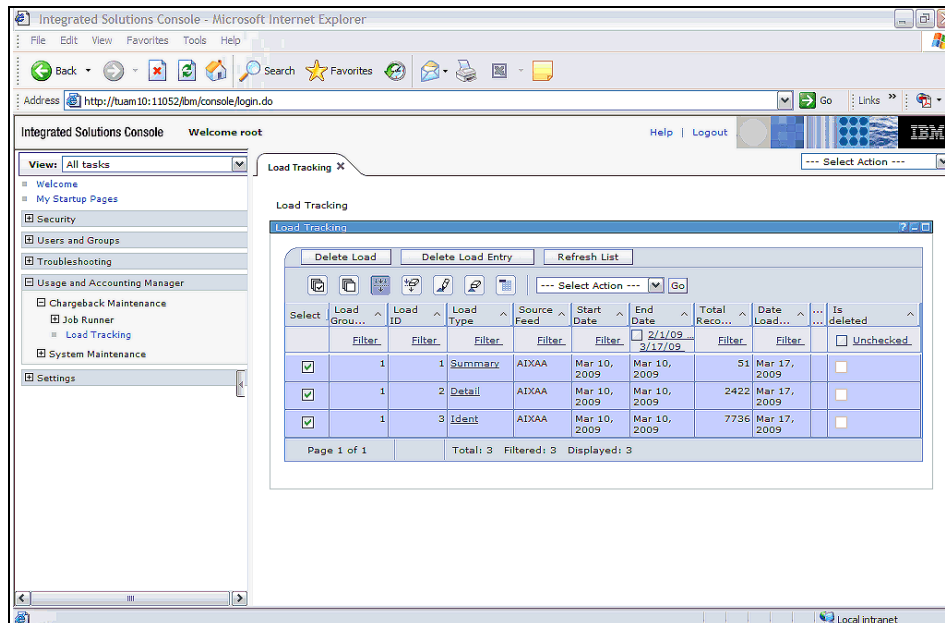
Figure 3-92   Load tracking

2. Click **Delete Load** and confirm the security question to delete the data.

3.5.7 AIX and Virtual I/O Server data collectors

Tivoli Usage and Accounting Manager Virtualization for System p includes collectors for:

► AIX or Linux Operating System (pacct)
► AIX or Linux File System
► AIX Advanced Accounting
► Virtual I/O Server data collector

The configuration of the data collectors must be automated in order to gather the data from all the systems in the environment back to the Tivoli Usage and Accounting Manager Server. You will need to install the data and configure the data collectors on each of the AIX or Virtual I/O Server systems where you want to gather data. The process of gathering AIX Advanced Accounting Data for use with Tivoli Usage and Accounting Manager is shown in Figure 3-93.
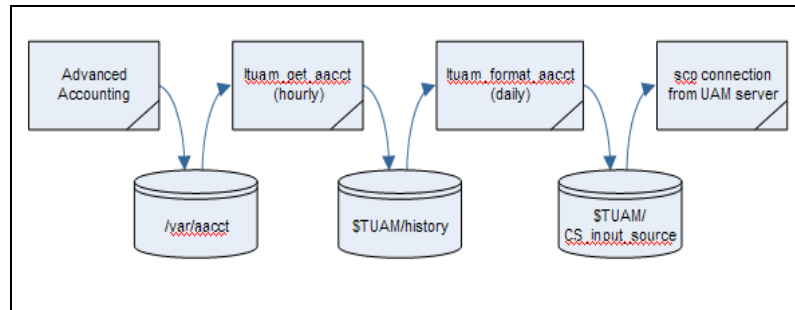
Figure 3-93   IBM Tivoli Usage and Accounting Manager data flow

3.5.8 AIX Advanced Accounting overview

AIX Advanced Accounting was introduced in AIX 5L™ V5.3, to provide more detailed accounting than the traditional UNIX® accounting. Advanced Accounting supports the new virtualization technologies that are available with Power Systems and AIX, such as logical partitions (LPARs), micro-partitioning, and workload partitions (WPARs). You can also configure and use AIX Advanced Accounting on the Virtual I/O Server.

The data sources for Advanced Accounting include resources, such as disks, network interfaces, virtual devices, file systems, processors, and memory. We discuss the eleven major transaction record IDs in Advanced Accounting. When you turn on accounting, you need to select to enable any of the transactions that you want to track. Table 3-4 on page 145 lists the transaction records along with the type of data that it will record.

Table 3-4   AIX Advanced Accounting record types

AACCT Transaction Record ID	Type of data
1-2	Process activity
4	System interval
6	File system activity
7	Network interface I/O
8	Disk I/O
10	Virtual I/O Server I/O
11	Virtual I/O Server client I/O
16	ARM-enabled transactions
36	WPAR system interval
38	WPAR file system activity
39	WPAR disk I/O

You can obtain more details about AIX Advanced Accounting in Chapter 5
"Advanced Accounting," in *Accounting and Auditing on AIX 5L*, SG24-6396.

3.5.9 AIX data collector installation and configuration

In our implementation, data collection on the client 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 using the `scp` command.

In Figure 3-93 on page 144, `$TUAM` represents our Tivoli Usage and Accounting
Manager data collector installation directory `/opt/ibm/tuam/collectors/unix`.

Installation prerequisites

For the installation of the AIX Advanced Accounting and Tivoli Usage and Accounting Manager data collector, we created a directory on the NFS server that contains the files that are described in Table 3-5.

Table 3-5   Tivoli Usage and Accounting Manager AIX collector installation files

File name	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_rsa.pub	SSH public key for root user of our Tivoli Usage and Accounting Manager server, which is required for setting up data collection and generated with the **ssh-keygen** command

You must also install the openssh packages on all the systems to which you will be adding data collectors. You can locate the openssh packages on the Expansion Pack CD-ROM that is included with AIX V6.1 Enterprise Edition.

Example 3-22 shows the installation script for IBM Tivoli Usage and Accounting Manager AIX collector.

Example 3-22   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 agram 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_rsa.pub ~${TUAMUSER}/.ssh/authorized_keys2
chown -R ${TUAMUSER}.${TUAMUSER} ~${TUAMUSER}/.ssh
chmod 755 ~${TUAMUSER} ~${TUAMUSER}/.ssh

#Set up the script for the 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_aix5.tar ${INSTPATH}
cd ${INSTPATH}
#use aacct_config=true when calling tuam_unpack_uc_collector to prevent
#traditional UNIX process accounting for being started
./tuam_unpack_uc_collector path=${INSTPATH} user=${TUAMUSER} aacct_config=true
rm ${INSTPATH}/tuam_unpack_uc_collector ${INSTPATH}/ituam_uc_aix5.tar
chmod 755 ${INSTPATH}/tuam_unpack_uc_collector ${INSTPATH}/ituam_uc_aix5.tar
#Adjust TUAM AAA collector to include client vscsi
cp ${INSTPATH}/data/A_config.par ${INSTPATH}/data/A_config.par.back
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

#Adjust the 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 3-23 on page 148 shows an example of the collector installation script output from one of our systems.

Example 3-23   The install-tuam.sh output

```
taddm10:/mnt/tuam ] ls -l
total 42248
-rwxr-xr-x    1 root     system          1921 Mar 19 15:18 install-tuam.sh
-rwxr-xr-x    1 root     system           450 Mar 19 11:43 ituam-schedule.sh
-rwxr-xr-x    1 root     system      21575680 Mar 19 11:45 ituam_uc_aix5.tar
-rwxr-xr-x    1 root     system           393 Mar 19 11:52 tuam10.pub
-rwxr-xr-x    1 root     system         36719 Mar 19 11:45 tuam_unpack_uc_collector

taddm10:/mnt/tuam ] ./install-tuam.sh
chmod: /opt/ibm/tuam/collectors/unix/tuam_unpack_collector: No such file or directory
./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                      : taddm10
    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                  : TRUE

./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_stop_odb, 10526 bytes, 21 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 - AACCT_ONLY
        Setting - ITUAM_SAMPLE
        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 Thu Mar 19 15:37:26 CDT 2009
*********************************************************************************


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 Thu Mar 19 15:37:34 CDT 2009
*********************************************************************************


./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
    previous ituam_uc_nightly entry replaced in crontab
./tuam_unpack_uc_collector: ITUAM UNIX/Linux Data Collector installation complete
```

Virtual I/O Server data collector set up

Tivoli Usage and Accounting Manager Virtualization Edition also supports usage reporting for Virtual I/O Servers. Because the Virtual I/O Server is a closed architecture, the data collectors come installed with Virtual I/O Server Version 1.4. We recommend fix pack 8.1.0 or later to run Tivoli Usage and Accounting Manager agents. Example 3-23 on page 148 shows the steps to turning on data collection on Virtual I/O Server. The commands are run by the Virtual I/O Server user padmin. The high-level steps are:

1. List all the available IBM Tivoli Usage and Accounting Manager agents using the **lssvc** command.

2. Using the **cfgsvc** command, list the attributes that are associated with the IBM Tivoli Usage and Accounting Manager agent that you want to configure.

3. Configure the IBM Tivoli Usage and Accounting Manager agent with its associated attributes using the **cfgsvc** command.

4. Start the IBM Tivoli Usage and Accounting Manager agent using the `startsvc` command.

Example 3-24   Configuration and start of data collector on Virtual I/O Server

```
$ lssvc
ITM_base
ITM_premium
TSM_base
ITUAM_base
TPC
DIRECTOR_agent
$ cfgsvc -ls ITUAM_base
ACCT_DATA0
ACCT_DATA1
ISYSTEM
IPROCESS
$ cfgsvc ITUAM_base -attr ACCT_DATA0=10 ACCT_DATA1=10 ISYSTEM=5
IPROCESS=5
$ startsvc ITUAM_base
```

3.6 Installing and configuring WPAR manager

In this section, we briefly describe the concept of workload partitions (WPARs) and the functions of WPAR Manager.

Note: For installing and configuring a WPAR Manager, refer to the following publication that describes WPAR and WPAR Manager in-depth: *Workload Partition Management in IBM AIX Version 6.1*, SG24-7656.

3.6.1 WPAR concept

In AIX 6, workload partitions add an additional operating system software-based layer for the 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 Version 6, including IBM pSeries POWER4 technology-based servers and later. Figure 3-94 on page 151 shows three application-specific or service-specific WPARs being hosted within a single LPAR.
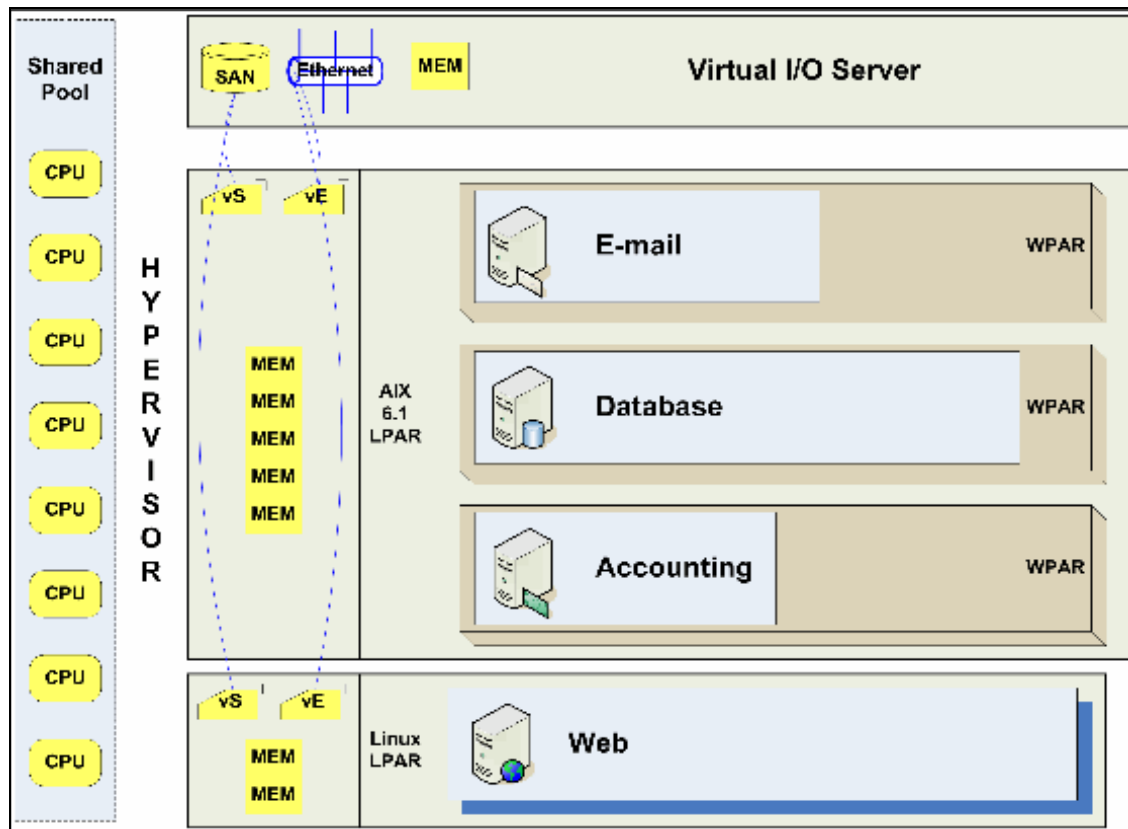
Figure 3-94   WPARs reduce LPARs but still provide workload isolation

Workload partitions can be created within an AIX V6 LPAR. Each workload partition provides an isolated environment for the application that it hosts. From the application or service point of view, the WPAR provides a 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

► 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

IBM Workload Partition Manager for AIX (WPAR Manager) is a tool for monitoring and managing WPARs. It will also gather information about CPU, memory, and other useful statistics.

There are new concepts introduced in WPAR:

► Global environment

This term refers to the part of the AIX operating system that hosts workload partitions. This environment is the classical AIX environment. Typically, only the AIX V6 system's root superuser has access to to this environment, because it must be set up to host WPARs exclusively, not native applications.

► System WPAR

This term refers to a more flexible WPAR-based instance of AIX. It contains dedicated writable file systems and system service daemons.

► Application WPAR

This term refers to a WPAR that is set up to host only a single application or process. It provides an AIX runtime environment that is suitable for execution of one or more processes that can be started from a single command.

3.6.2 WPAR Manager components and functions

WPAR Manager is a Web-based Java application that uses Common Agent Services (CAS) and a database. CAS is installed automatically when you install WPAR Manager.

WPAR Manager can use either Apache Derby or DB2 as a database. During WPAR Manager installation, you select which database to use. If DB2 is selected, you also have to install and configure DB2, which is a separate installation process.

After the basic installation, CAS, WPAR Manager, and the database must be configured. When these configurations are complete, WPAR Manager is ready to orchestrate mobility events among servers.

Common Agent Service

To help you understand how WPAR Manager components work, we provide a brief overview of the Common Agent Services (CAS) framework. Figure 3-95 on page 153 shows WPAR Manager components, including the management server and two managed servers.

In this scenario:

► WPAR Manager is the application and user interface. It is in charge of communicating with the user and updating the database.

► CAS Agent Manager is in charge of the communication between the WPAR agent and the WPAR Manager in the LPAR.

► The database is the repository for all WPAR definition data.

There are WPARs running on the managed servers. B and G are set up as system WPARs, and W is set up as an application WPAR.

In this setup, the WPAR Manager must be able to move WPARs W and G between the two nodes.
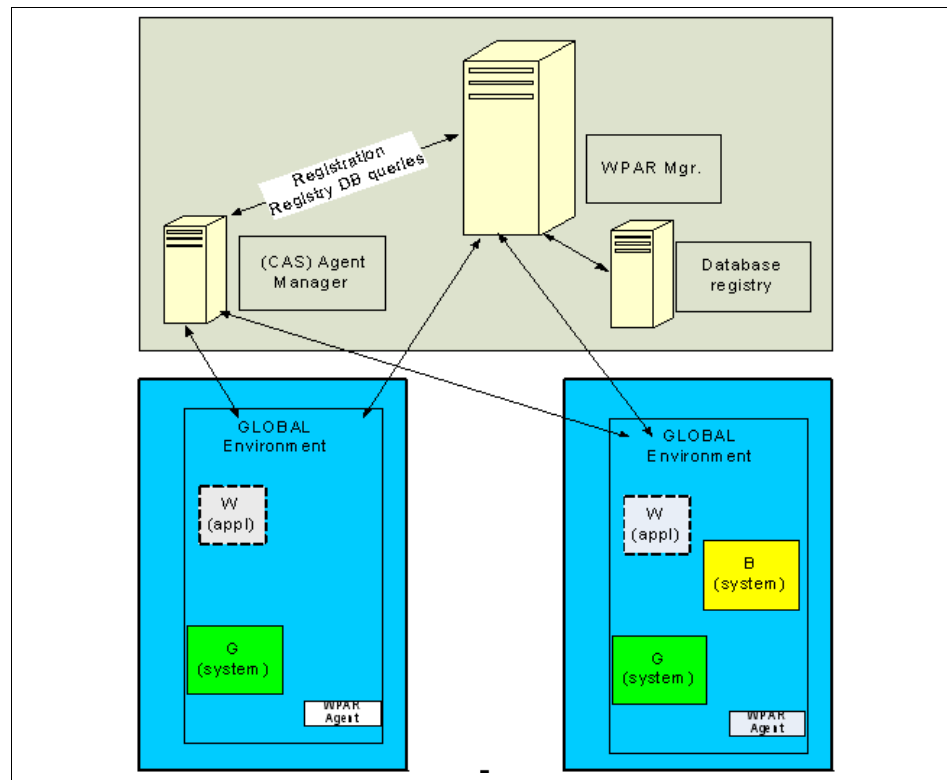
Figure 3-95   Components of WPAR Manager

Deploying WPAR Management software usually requires a server that hosts the management software and an agent that has to be installed on each server to be managed, as shown in Figure 3-95.

The effort required to install, configure, maintain, and monitor various agent implementations can become an extremely time-consuming management task. In addition, the number of redundant functions that is provided by each agent implementation (for example, listening to ports, the runtime daemon, and memory consumption) leads to ineffective usage of the system resources.

The goal of Common Agent Services is to minimize the complexity of the software management deployment by reducing the effort that is needed for deployment and utilizing system resources more effectively.

CAS accomplishes this goal by providing a common framework that can be reused by various management applications that are deployed across the enterprise.

This Common Agent Services framework consists of the following components:

► Agent Manager
► Resource Manager
► Common Agent

Agent Manager

Note: The Agent Manager component implements the CAS in the WPAR Manager application. This agent talks to the Common Agent on the managed systems.

The *Agent Manager* is the server component of the Common Agent Services. It provides authentication and authorization services, enables secure connections between managed systems in your deployment, and maintains the registry about the managed systems and the software running on those systems. It also handles queries from the Resource Managers against the database.

The Agent Manager has the following components:

► The Agent Manager service

   The Agent Manager service serves as a certificate and registration authority to provide authentication and authorization using X.590 digital certificates and the Secure Sockets Layer (SSL) protocol. It also handles requests for registry information from Common Agents and Resource Managers.

   Resource Managers and Common Agents must register with the Agent Manager service before they can use its services to communicate with each other. This registration is password-protected, and there are separate passwords for the Common Agents and Resource Managers.

> **Note:** For WPAR Manager, you only need to specify the registration password for the Common Agents. The password for the Resource Manager is automatically generated during the configuration of WPAR Manager.

► The registry

The registry is the database that contains the current configuration of all known Common Agents and Resource Managers.

Information contained in the registry includes:

– The identity, digital certificates, and communication information for each Resource Manager

– The identity, digital certificates, and communication information for each Common Agent

– Basic configuration information for each Common Agent (for example, hardware type and operating system version)

– The status of each Common Agent

– The last error or, optionally, a considerable number of errors, reported by each Common Agent

The registry information is updated by events, such as the registration of Common Agents and Resource Managers, and by periodic updates from the Common Agents.

> **Note:** For WPAR Manager, the Agent Manager uses an Apache Derby database to implement the registry. This Derby database is included within the lightweight runtime fileset that is delivered with WPAR Manager.

► The agent recovery service

The agent recovery service provides error logging for Common Agents that cannot communicate with other Agent Manager services.

WPAR Manager does not make use of this feature.

> **Note:** The Agent Manager functionality for WPAR Manager is in the fileset `wparmgt.cas.agentmgr`.
>
> CAS Agent Manager listens to communication from WPAR Agent and WPAR Manager on ports 9511, 9512, and 9513. These ports are default ports that can be overridden by the user during the configuration of WPAR Manager.

Resource Manager

A *Resource Manager* is management software that uses the services of the Agent Manager to communicate securely with, and obtain information about, the managed systems running the Common Agent software. The Resource Manager uses the services of the Common Agent to deploy and run its software on the managed systems.

Note that in our case, the WPAR Manager component of WPAR Manager is the Resource Manager of the CAS framework.

> **Note:** The Resource Manager functionality for WPAR Manager is in the fileset `wparmgt.mgr.rte`. This fileset also includes the functionality of the WPAR Management console.
>
> WPAR Manager listens to ports 14080 and 14443 and communicates to port 9510 on the WPAR Agent. WPAR Manager listens to ports 9511, 9512, and 9513 on the CAS Agent Manager. These ports are default ports that can be overridden by the user during the configuration of WPAR Manager.

Common Agent

Each managed system in your deployment runs a *Common Agent*, which provides a single implementation of common services that can be used by all management products.

The Common Agent consists of:

► Common Agent services

► Product-specific subagent (for example, for WPAR Manager, this code is the code for WPAR Agent)

For instance, if two management products (A and B) manage the same system, the Common Agent is installed only one time on that system, and it runs two product-specific subagents: one for product A and one for product B.

The Common Agent provides the following functions:

► A single set of security credentials and a common security infrastructure for all management applications

► Automated management of security credentials

The Common Agent certificate is automatically renewed when it is near the expiration date.

► Common Agent health monitoring and configuration monitoring services

The Common Agent has a "heartbeat" function that sends periodic status and configuration reports to the Agent Manager. It also allows any subagent to participate and to provide application-specific status.

The frequency of this update can be set, or it can be turned off. Management applications can register to the Agent Manager to receive these updates.

The registry contains only the most recent configuration update. By default, only the most recent status information is saved for each Common Agent. The retention period is configurable.

Operations

This chapter describes and discusses several of the more common optimization and utilization features of the components in day to day operations within a Power System environment:

4.1 Live Application Mobility

The ability to relocate a system's or an application's workload partitions (WPARs) from one AIX host or global environment to another AIX host or global environment means that we can optimize our virtualized System p hosts by deploying and redeploying WPARs using application mobility manually or automatically based on policies.

Policy-based Live Application Mobility assists in the overall utilization of WPARs within an environment by balancing the workloads across a range of Power Systems within the environment without administrator intervention.

Note: Live Application Mobility currently relies on common Network FIle System (NFS)-mounted file systems between the two global environments or AIX images. Future support for SAN and GPFS file systems is expected.

There are two modes of operation available for relocation: a classic mode that stops execution within the WPAR during the migration and an asynchronous mode that allows execution to continue during the migration. During the mobility operations, the application undergoes active relocation (*hot migration*) without stopping or disrupting the application.

Note: Live Application Mobility is not designed to, in any way, be a replacement for a High Availability (HA) solution. However, live application mobility can complement the IBM PowerHA offering.

Application mobility provides us with not only another optimization tool to better utilize our physical resources, it also allows for better management of planned migrations of workloads from one system to another system during planned outages, for instance, such as hardware maintenance, firmware installation, or energy conservation. The workload does not need to be aware of the migration, but we always recommend that you perform proper planning and testing before moving anything into or around a production environment.

For further information, refer to:

http://www.redbooks.ibm.com/abstracts/sg247656.html?Open

4.2 Operating TM

The following sections outline the basic operations involved with IBM Tivoli Monitoring:

4.2.1 Workspaces

The IBM Tivoli Monitoring tool that is packaged as part of AIX Enterprise Edition comes bundled with a range of queries and views to assist you with getting the tool operational in a short amount of time. You can then create your own views and queries, based on the predefined offerings for AIX, Virtual I/O Server, Hardware Management Console (HMC), and common event console (CEC).

4.2.2 Views

A range of predefined workspaces are built into your IBM Tivoli Monitoring product that consists primarily of chart views and table views. You can save them as new workspaces and customize them with a range of views.

Data views

The table and chart views are the first steps to getting something meaningful from the data that is being collected. When you understand what values and states cause problems, you can refine your views to show what is important.

The table view and the chart view provide current and historical information about monitoring data from Tivoli Enterprise Monitoring agents, such as system performance and configuration:

► Pie charts show portions of a whole (such as a percentage).

► Bar and plot charts, as well as gauges, show the values of each attribute.
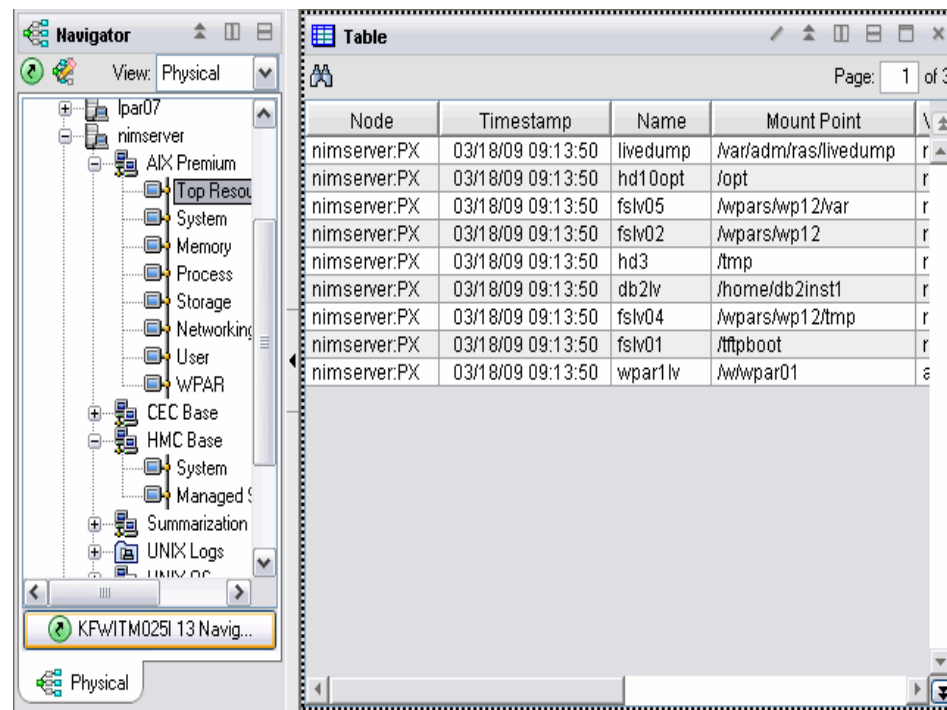
The following sections provide an explanation of each of the views and provide an example of the same system metrics from our Network Installation Management (NIM) server represented in the various views. Each view can be

initiated by clicking the corresponding icon from the Tivoli Enterprise Portal Toolbar and dragging the icon onto the panel that you want displayed in that particular view.

Table views

Table views show current data that is retrieved from a monitoring agent and show one column for every attribute. The table reports a single row of data or multiple rows, depending on the nature of the attribute group.

The table view also has features for refining the display to suit your needs. These features include adjusting column width, locking columns, and using drag and drop to change the order of columns. You can also sort column data differently in ascending or descending sequence by clicking the column heading. Figure 4-1 shows a table view.

Figure 4-1   Table view

Chart views

You can view the predefined workspaces in one of five ways as shown in
Table 4-1 on page 168:

► A *pie chart* shows a slice for every data point in a single data series (row). Pie
charts are well suited for showing the proportional value of related attributes
to a whole, such as the percentage attributes that show how memory is being
used (Figure 4-2).
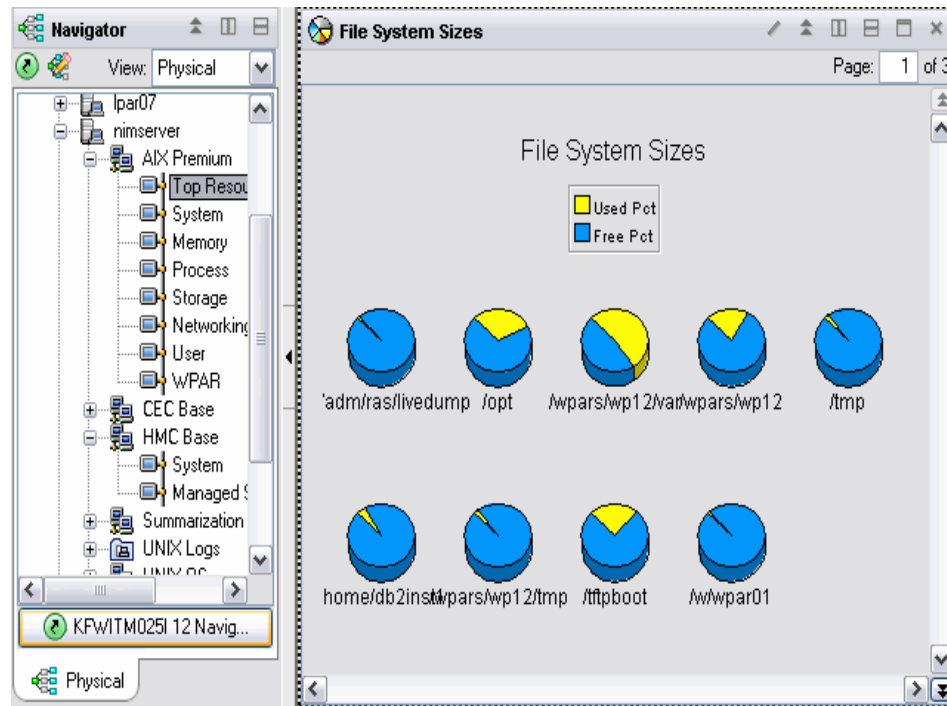
Figure 4-2   Pie chart

► A *bar chart* displays a bar for each data point. Bar charts are best suited for comparing values among related attributes. The stacking bar chart is well suited for showing multiple values for the same attribute (Figure 4-3).
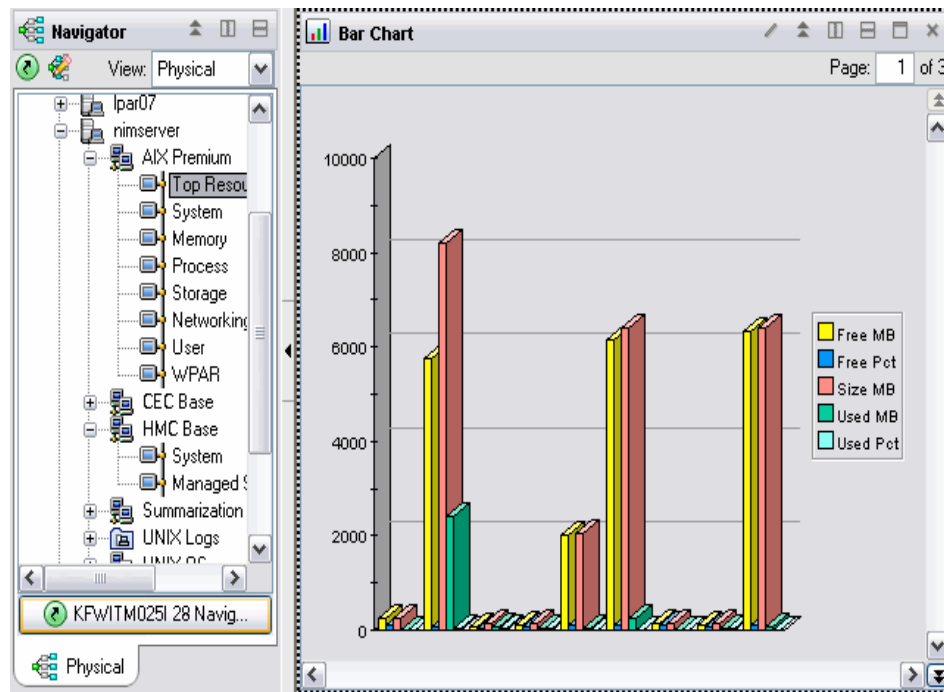
Figure 4-3   Bar chart

► A *plot chart* shows changes over a period of time by drawing a continuous line from one data point to the next data point with one data point for each data sampling and one line for each selected attribute. The plot chart can be used for plotting multiple-row attribute groups (or historical data from a single-row attribute group) and multiple managed systems. You can also control the refresh rate of the plot chart so that it is independent of the refresh rate of the workspace as a whole. The plot chart is well suited for showing trends over time and among related attributes (Figure 4-4).
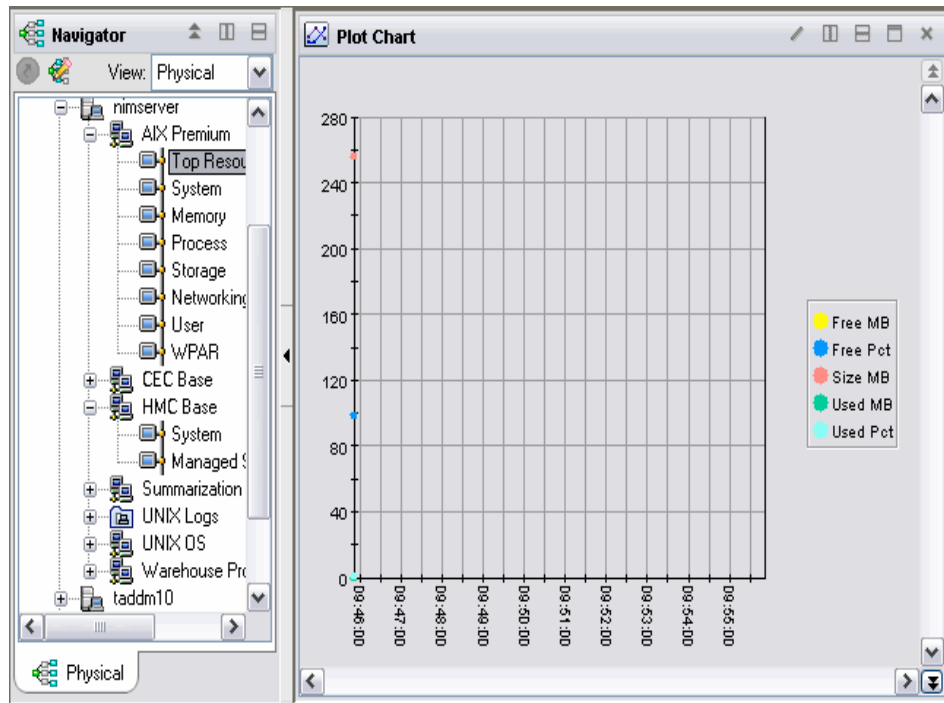
Figure 4-4   Plot chart

► A *circular gauge* shows the proportional amount of a data series with one gauge for each chosen attribute. This type of chart is well suited for showing individual elements that change frequently, such as a percentage of user time (Figure 4-5).
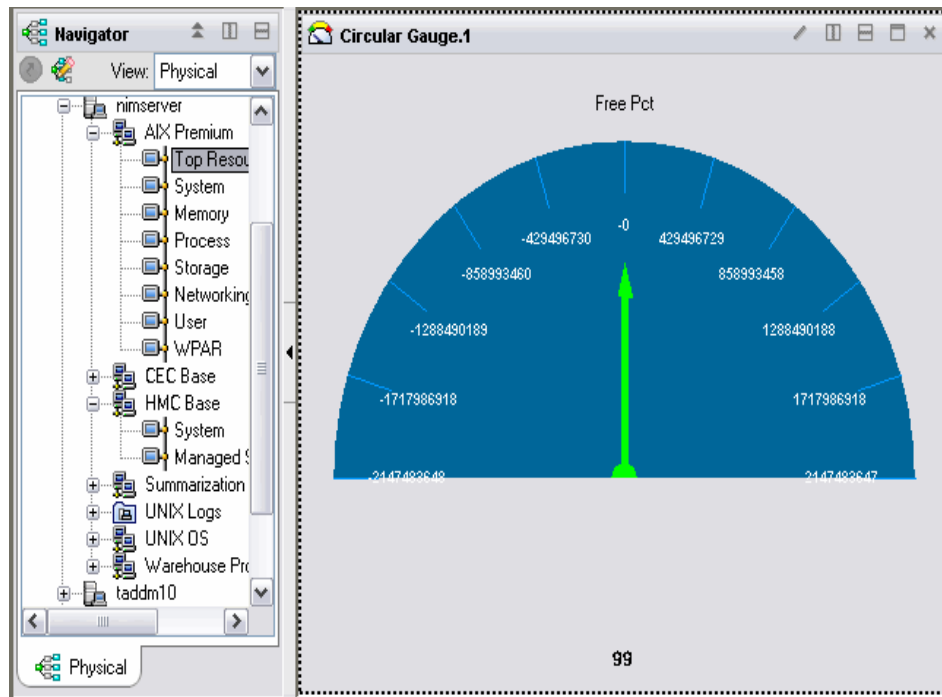
Figure 4-5   Circular gauge

► A *linear gauge* shows the collective value of every item in a single data series with one gauge for each chosen attribute. This type of chart is well suited for showing cumulative values (Figure 4-6).

Figure 4-6   Linear gauge

In Table 4-1 on page 168, we look at the best rendering of the individual views based on the data to be displayed. This view changes from query to query and largely depends on how you configure the workspace.

Table 4-1   Chart types and their best uses
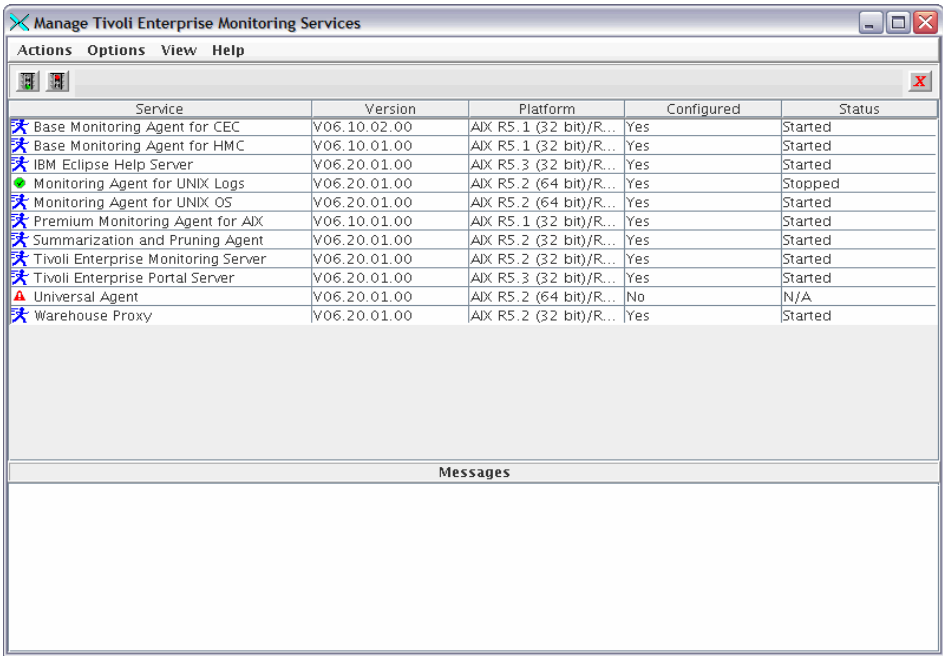
Chart type	Best for	Multiple attributes	Multiple rows	Multiple managed systems	Time span
Pie chart	Displays proportional value to the whole	One slice per attribute	One pie for each row	One pie for each managed system	Yes
Bar chart	Compares values among related attributes	One bar per attribute; stacking bars show one segment per attribute	No	No	Yes
Plot chart	Displays trends over a period of time for related attributes	One line per attribute; one data point for each data sampling	Yes	Yes	Yes
Circular gauge	Shows frequently changing elements	One gauge per attribute	No	No	No
Linear gauge	Displays cumulative values	One gauge per attribute	No	No	No

4.2.3 Managing Tivoli Monitoring Services

Manage Tivoli Enterprise Monitoring Services, which is seen in Figure 4-7 on page 169, is often the component of the IBM Tivoli Monitoring V6.2 product suite that is overlooked. The Manage Tivoli Enterprise Monitoring Services tools functionality includes:

► Install product support (TEMS)
► Add application support (TEMS)
► Start and stop the server or agents
► Make configuration changes
► Manage log files

You can launch the Manage Tivoli Enterprise Monitoring Services application by exporting your display, ensuring that you have an X Windows Server, such as Hummingbird's Exceed, running, and entering the command `/opt/IBM/ITM/bin/CandleManage` or `$CANDLEHOME/bin/CandleManage` on any of the machines that have agents or server components installed. Refer to Figure 4-7.

Figure 4-7   Manage Tivoli Enterprise Monitoring Services

4.2.4 Historical data collection

The ability to collect and store historical data allows us to examine the performance of the monitored components over a period of time and to provide a sound basis for forecasting growth and capacity planning, as well as allowing us to manage our virtualized components within the environment to maximize efficiencies.

How it works

If historical data has been configured, it is collected by either the monitoring agent or the TEMS in the short term. The data is stored in tables, with one table per attribute that is collected. If the Tivoli data warehousing facility has been set up, the collection goes into either an Open Database Connectivity (ODBC) or Java Database Connectivity (JDBC) database for long-term storage.

If historical data has been enabled for the attributes for a view, there is a time span icon on the toolbar. To view the historical data instead of the real-time feed in the workspace view, you use the tool to specify the period and duration for the data that you want to display. When the query is run, data less than 24 hours old is extracted from the sourced binary files, while the data older than 24 hours is taken from the data warehouse.

The History Collection Configuration window in the Tivoli Enterprise Portal enables you turn the collection on or off and to configure how often and where the historical data is stored.

> **Note:** Data collection cannot be started unless you have configured it.

Configuring data collection and warehousing

Take the following steps to configure historical data collection:

1. Click the 🔲 History Configuration icon from the Edit menu in the TEPS to open the History Collection Configuration window that is shown in Figure 4-8.
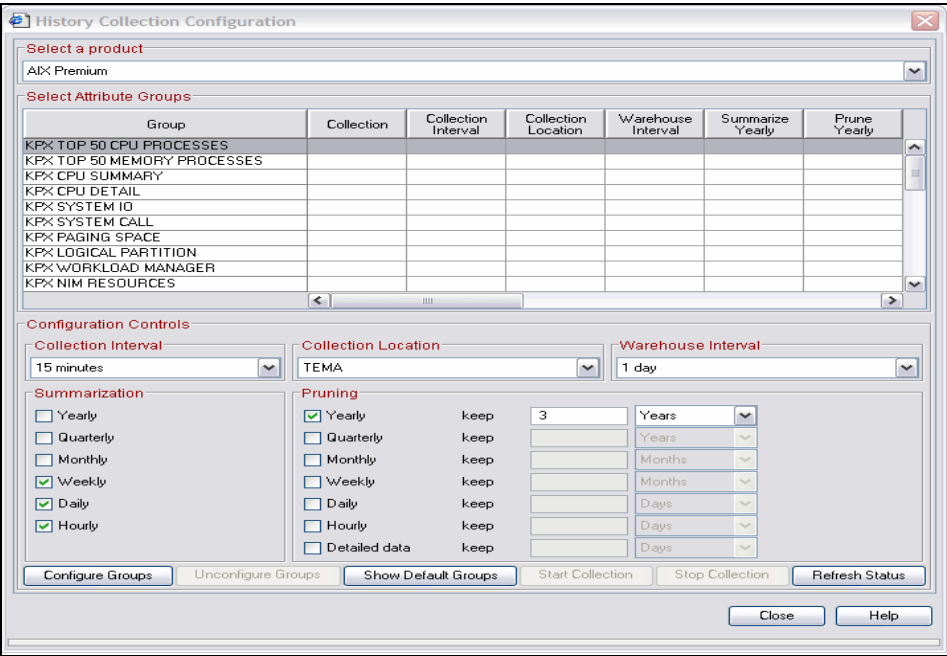
Figure 4-8   History Collection Configuration window

2. Select the agent type for which you want to collect data from the **Select a product** drop-down list.

   The attribute groups for which you can collect historical data are displayed in the **Select Attribute Groups** list box. Note that when you select a product type, you are configuring collection for all monitoring agents of that type that report to the selected monitoring server. If your monitored environment is large scale with many hundreds of agents of a type, there might be a delay when you select the product.

3. Select one or more attribute groups. You can also click **Show Default Groups** to automatically select the attribute groups that have predefined historical workspaces.

   If you have configurations that you want to populate to several attribute groups, click the attribute group that has the configurations that you want to see reflected in the Configuration Controls area, then highlight the other attribute groups (Ctrl+click to select multiples, or Shift+click to select all groups from the first one selected up to this point) to configure and click **Configure Group**. All the selected attribute groups will reflect the same configurations. If collection has been started for any of the attribute groups that you have selected, stop the collection before attempting to change the configurations.

4. Specify the following collection options:

   – **Collection Interval**

     Specify the frequency of data sample transmission to the history file. The options are: every 1, 5, 15, or 30 minutes, every hour, or every day. The default interval is 15 minutes. The shorter the interval, the faster and larger the history file grows at the collection location. This history file can overload the database, Warehouse Proxy agent, and Summarization and Pruning agent. For example, if you set a 1 minute collection interval for process data, expect the summarization and pruning for that attribute group to take a long time. Only enable such a short interval for an attribute group that is critical in your work and then preferably only for a short period of time.

     Note that selecting too short of an interval can quickly fill file systems. We recommend that you monitor your data collection rates following any changes to the collection settings.

   – **Collection Location**

     Specify where the historical data files will reside: at the Tivoli Enterprise Monitoring Agent (TEMA) or the Tivoli Enterprise Monitoring Server (TEMS). The default location is the monitoring agent, which minimizes the performance impact on the monitoring server from the management of

historical data; however, the TEMS might be a better choice for certain environments.

– **Warehouse Interval**

Specify whether the collected data is warehoused and how often. The options are 15 minutes, 30 minutes, 1 hour, 12 hours, 1 day, or Off. If you choose to warehouse historical data, the scheduling options for Summarization and Pruning are enabled. A more frequent warehousing interval enables quicker availability of warehoused data for retrieval. Shorter intervals cause additional processing to check for and transmit the newly collected data. And there are more frequent elevated levels of transmission activity, but for shorter durations.

The collection options for an attribute group apply on all monitoring servers on which collection for those attributes is currently enabled. For certain sets of attribute groups, the configuration is hard-coded by the product. These attribute groups will display a (`fixed`) in the Collection and Warehouse Interval cells of the Select Attribute Groups table.

5. If warehousing is enabled, specify the time periods for which data in the warehouse is summarized.

6. If summarization is scheduled, specify the duration to keep the summarized data for each time period.

7. If pruning is scheduled, specify how long to keep the pruned data for each time period.

8. Click **Configure Groups** to apply the configuration selections to the attribute group or groups. If you have currently configured controls for this group (and the controls are running), click **Stop Collection** first. Apply your new settings, and click **Start Collection**.

> **Note:** Your user ID must have permissions sufficient to edit workspaces in the Tivoli Enterprise Portal in order for you to be able to configure, start, or stop the historical data collection.

4.2.5 Understanding situation events

Situations notify you when an event occurs on a monitored system. The IBM Tivoli Monitoring product provides a set of predefined situations for your immediate use to expedite the time frame in which you can have your environment monitored. You can view and edit these situations and create other customized situations in the Situation Editor to suit your specific environment.

The Situation Editor window is divided into two frames. The left frame contains the Situation tree, which shows the situations associated with the Navigator item from which the Situation Editor was opened as shown in Figure 4-9, or the situations for all installed monitoring products if you opened the editor from the toolbar (the situations show in the left frame of Figure 4-9).
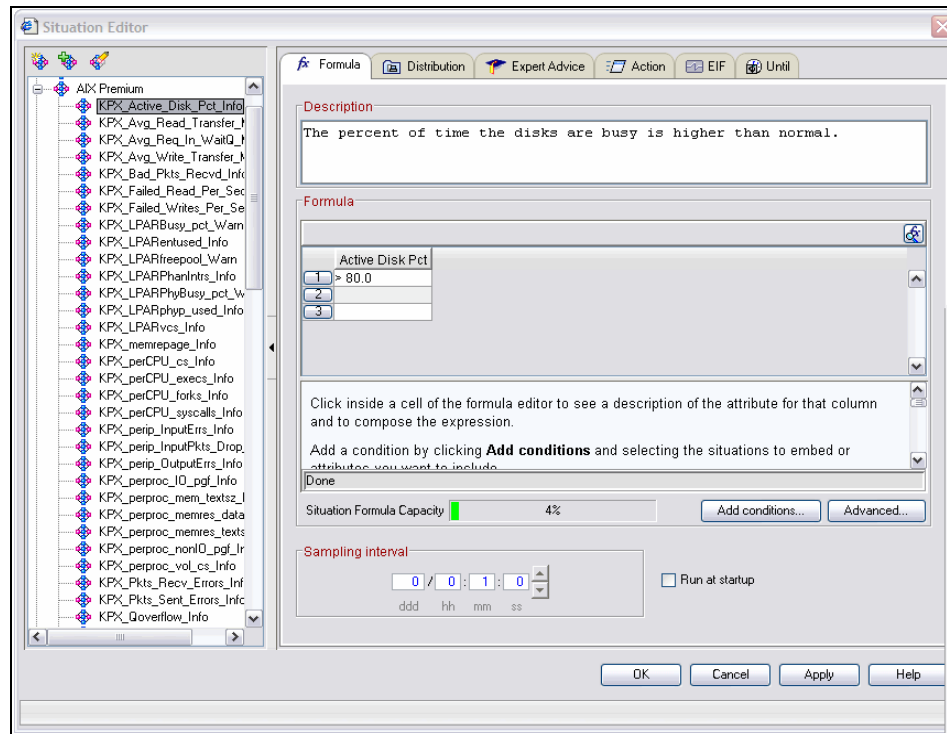
Figure 4-9   Situation Editor

Situation formula

Situation formulas are one or more expressions of which options are given when you click the Formula button in the Situation Editor window. For example, when a situation checks for free disk space below 10%, it uses the group and item attributes that we see in Figure 4-10 on page 174.

You can also create more elaborate situations that are based on multiple conditions being met before an event is triggered.
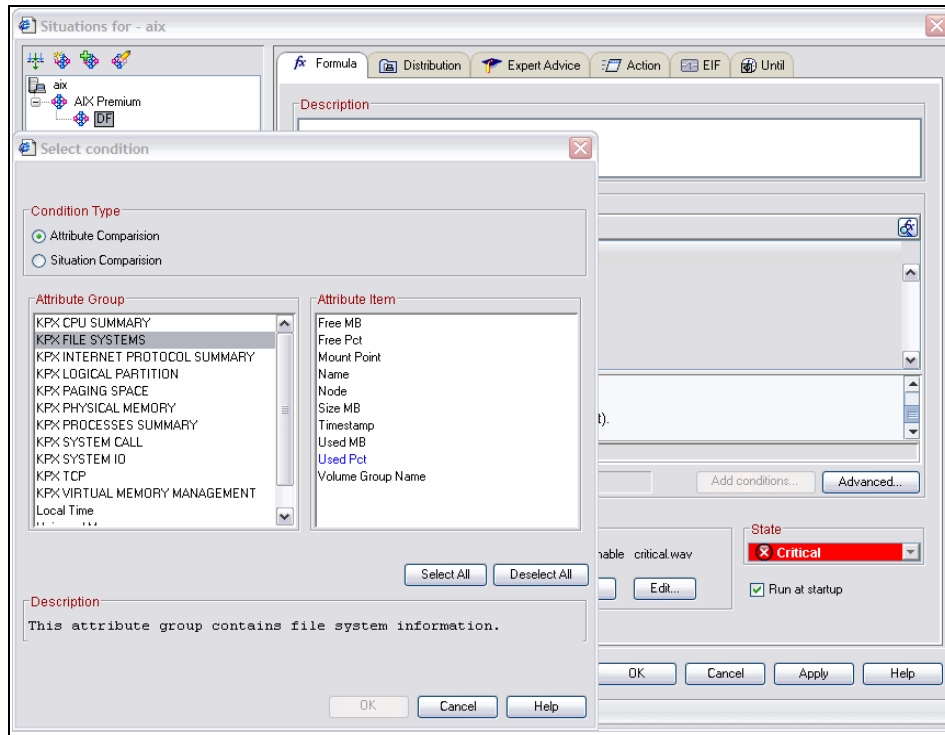
Figure 4-10   Attributes of a situation formula

Situation event indicators

When a situation triggers an event on a managed system, it also has a *state* setting, which determines the light and color displayed in the *Navigator* window on the item for which the situation threshold has been reached. The list of states running from highest to lowest severity can be seen in Figure 4-11.

Figure 4-11   Situation event indicators (states)

4.2.6 IBM Tivoli Monitoring command line

There are a few major commands in IBM Tivoli Monitoring V6.2 that you can use to modify, stop, start, and reconfigure the server or agents. The major commands are:

`tacmd`   Tivoli Agent Command

`itmcmd`   IBM Tivoli Monitoring command

`cinfo`   This command returns a menu to query information, such as installed and running agents

For further information about these commands, refer to *IBM Tivoli Monitoring User's Guide, Version 6.2.0,* SC32-9409, and *IBM Tivoli Monitoring Administrator's Guide, Version 6.2.0,* SC32-9408. This section focuses on providing the essential information about these commands.

Common commands
Basic `tacmd` options are:

► `help`

Prints the full help text for the specified command or summary help text for all commands if no command is specified.

► `login`

Authenticates a user with a server and creates a security token that is used by subsequent `tacmd` commands.

► `logout`

Invalidates the security token created by the tacmd `login` command.

► `refreshTECinfo`

Triggers the Event Forwarder to reprocess any updated event destinations, Event Integration Base (EIF) configurations, and custom event mapping files without recycling the hub Tivoli Enterprise Monitoring Server.

Installation commands
The `tacmd` installation commands are:

► `createNode`

Creates a node by installing an OS agent into a new directory on the local machine and starting that OS agent.

► `viewNode`

Displays the details of a node, including the components that are installed on it.

Depot commands

The depot `tacmd` commands are:

► **viewDepot**

Displays the types of agents that can be installed from the (currently logged in) server's agent depot or from the named remote server's depot.

► **listBundles**

Displays the details of one or more deployment bundles that are available for installation from the specified directory into the local deployment depot.

► **addBundles**

Installs one or more deployment bundles from the specified directory into the local deployment depot. This command must be executed locally on a Tivoli Enterprise Monitoring Server containing a depot.

► **removeBundles**

Removes one or more deployment bundles from the local deployment depot. This command must be executed locally on a Tivoli Enterprise Monitoring Server containing a depot.

Agent commands

The agent `tacmd` commands are:

► **startAgent**

The **startAgent** command:

– Starts the specific agents or the agents for the specific managed systems if they are not already running.

– OS agents (nodes) can be started on the local system only.

– The agent is marked online in the Tivoli Enterprise Portal client.

► **stopAgent**

The **stopAgent** command:

– Stops the specific agents or the agents for the specific managed systems if they are running.

– OS agents (nodes) can be stopped on the local system only.

– The agent is marked offline in the Tivoli Enterprise Portal client within the next heartbeat interval.

► **restartAgent**

The **restartAgent** command:

– Starts or restarts the specific agents or the agents for the specific managed systems.

– OS agents (nodes) can be restarted on the local system only.

– If the agent is already started, it is stopped before being restarted.

► **viewAgent**

The **viewAgent** command:

– Displays the details of the specific agent or the agent for a specific managed system.

– Details include the agent version, agent status (running or not), and all of the configuration data for the agent.

► **updateAgent**

Enables an agent update to a specified node.

System commands

The **tacmd** system commands are:

► **describeSystemType**

Displays the configuration options that are required for a specific managed system type.

► **addSystem**

The **addSystem** command:

– Enables a user to add managed systems to the monitoring system.

– Deploys an agent and other needed components if they are not already installed on the node.

► **configureSystem**

The **configureSystem** command:

– Enables the user to change the configuration options of an existing managed system.

– The user also has the option to restart the managed system's monitoring agent in order for the new configuration parameters to take effect.

► **listSystems**

The **listSystems** command:

– Displays a list of known managed systems.

– The results can optionally be filtered to display only those results on a specific node, only those results for given product types, or both.

Server command

The server **tacmd** command is **configurePortalServer**. The **configurePortalServer** command:

► Adds, configures, or removes a portal server data source from the portal server configuration.

► This command must be executed locally to the portal server.

► This command uses NAME=VALUE options.

Situation commands

The situation **tacmd** commands are:

► **createSit**

Creates a new situation, based on an existing situation.

► **editSit**

The **editSit** command:

– Edits an on-server or exported situation definition.

– If a –force parameter is specified, it disables the message that asks if you are sure that you want to edit the situation.

► **deleteSit**

The **deleteSit** command:

– Deletes a situation from the server.

– If a –force parameter is specified, it disables the message that asks if you are sure that you want to delete the situation.

► **listSit**

Lists the defined situations, optionally filtering for a specific managed system or for a specific managed system type.

► **viewSit**

Allows the configuration of a situation to be displayed or saved in an export file.

Event destination definition commands

The event destination definition `tacmd` commands are:

► **createEventDest**

Creates a new event destination definition on the server.

► **deleteEventDest**

Deletes an event destination server definition from the server.

► **editEventDest**

Modifies an existing event destination server definition on the server.

► **listEventDest**

Displays the server ID, name, and type for every event destination definition on the server.

► **viewEventDest**

Displays all properties for the specified event destination definition on the server.

Import and export commands

The import and export `tacmd` commands are:

► **bulkExportPcy**

Exports all the available policies from the Tivoli Enterprise Monitoring Server.

► **bulkExportSit**

Exports all the available situations from the Tivoli Enterprise Monitoring Server.

► **bulkImportPcy**

Imports all the available policy objects to the Tivoli Enterprise Monitoring Server from BULK_OBJECT_PATH.

► **bulkImportSit**

Imports all the available objects to the Tivoli Enterprise Monitoring Server from BULK_OBJECT_PATH.

Managed systems list commands

The managed system list `tacmd` commands are:

► **createsystemlist**

Creates a new managed system list.

- ► **`deletesystemlist`**

  Deletes a managed system list.

- ► **`editsystemlist`**

  Adds or deletes managed systems to or from an existing managed system list on the server.

- ► **`listsystemlist`**

  Lists the available managed system lists.

- ► **`viewsystemlist`**

  Lists the configuration of a managed system list to be displayed or saved in an export file.

Workspace commands

The workspace **`tacmd`** commands are:

- ► **`exportworkspaces`**

  Exports one or more portal server workspaces to a file.

- ► **`importworkspaces`**

  Imports the workspaces contained in a file into the portal server.

- ► **`listworkspaces`**

  Lists all of the portal workspaces on the server.

The itmcmd command

The IBM Tivoli Monitoring command or **`itmcmd`** has the following format:

`itmcmd command [option [operand ...]] ...`

The **`itmcmd`** commands are:

- ► **`itmcmd agent`**

  The **`itmcmd agent`** command:

  – Starts and stops most agents.

  – You can start or stop one agent, all agents, or multiple agents.

  – You can also start the portal server and portal desktop client using this command.

  – You must run the **`itmcmd agent`** command on the architecture for which the agent is installed.

► **`itmcmd audit`**

Manages log files from the command line. The default logs are in `/opt/IBM/ITM/logs`.

► **`itmcmd config`**

The **`itmcmd config`** command:

– Configures or reconfigures the following components:

- The IP port that the hub monitoring server uses to listen for requests
- The hosts that can run a product
- The location of the hub monitoring server in the network
- The monitoring server to which an agent connects
- Whether a monitoring server is a hub or a remote server

– You can only configure one product at a time. If you reconfigure a monitoring server, you must stop and restart that monitoring server before the changes take effect.

► **`itmcmd dbagent`**

Starts the IBM Tivoli Monitoring for Sybase and IBM Tivoli Monitoring for Oracle monitoring agents.

► **`itmcmd dbconfig`**

Configures the execution environment for a distributed database agent.

► **`itmcmd execute`**

The **`itmcmd execute`** command:

– Runs a script or command when its execution requires the same environment settings as a particular IBM Tivoli product. The **`itmcmd execute`** command runs this script or command by building the necessary environment settings for the intended script or command and then combining them into a temporary shell script before running it.

– You must run the **`itmcmd execute`** command. Make sure that you are in the correct directory: **`cd $CANDLEHOME/bin`**, where **`$CANDLEHOME`** is the location where you installed your IBM Tivoli software, generally */opt/IBM/ITM*.

► **`itmcmd history`**

Manages the roll-off of history data into delimited text files.

► **`itmcmd manage`**

Starts Manage Tivoli Enterprise Monitoring Services on a UNIX or Linux computer. You can start, stop, and configure monitoring components in Manage Tivoli Enterprise Monitoring Services.

- **itmcmd server**

  Starts and stops monitoring servers that are defined in directories under the `install_dir/tables` subdirectory. You must run the **itmcmd server** command from the host computer.

- **itmcmd support**

  The **itmcmd support** command:

  – Adds agent-specific information to the monitoring server.

  – Whenever you add a new monitoring agent type to your monitoring environment, run the **itmcmd support** command again on the monitoring server to add the new agent information to the monitoring server.

4.2.7 The cinfo command

An administrator can utilize the **cinfo** commands to complete any of the following tasks:

- An inventory of installed IBM Tivoli products

- The configuration settings for products

- The installed CD versions in the current `$candlehome` directory

- The configuration settings for products in the context of the actual variables that are used by the installation program

- A list of running IBM Tivoli processes (such as agents or monitoring server)

- A validated list of running IBM Tivoli processes, after first performing an update on the tracking database to remove stale PIDs (processes logged as ps but not found when attempting to verify using the **ps** command)

This command can be menu-driven (run with no flags) or controlled by passing arguments as shown in Example 4-1 on page 183.

Example 4-1   The cinfo command

```
cinfo [-h candle_directory] [-c product] [-i] [-r] [-s product] [-R]
[-v]
   -c <product> Displays configuration prompts and values
   -i Displays an inventory of installed products
   -r Shows running processes
   -s <product> Displays configuration parameters and settings
   -R Shows running processes, after updating a tracking database
   -v Shows the installed CD versions in this CandleHome
   -p <product> Shows associated platform codes for the specified
   product
   -d Dumps an inventory of installed products
```

For example, if you run **cinfo -p ms**, you get information about the Tivoli
Enterprise Monitoring Server (TEMS) as shown in Example 4-2.

Example 4-2   The cinfo product details for the TEMS

```
root@nimserver:/opt/IBM/ITM/bin# cinfo -p ms

*********** Wed Mar 18 16:49:41 CDT 2009 ******************
User: root Groups: system bin sys security cron audit lp
Host name : nimserver    Installer Lvl:06.20.01.00
CandleHome: /opt/IBM/ITM
***********************************************************
Platform codes:
aix536 : Current machine
aix523 : Product (ms)
tmaitm6/aix523 : CT Framework (ax)
```

You can view the Tivoli Enterprise Monitoring processes that are running by
using **cinfo -r** as shown in Example 4-3 on page 184.

Example 4-3   Running Tivoli Enterprise Monitoring processes

```
root@nimserver:/opt/IBM/ITM/bin# cinfo -r

*********** Wed Mar 18 16:57:58 CDT 2009 ******************
User: root Groups: system bin sys security cron audit lp
Host name : nimserver    Installer Lvl:06.20.01.00
CandleHome: /opt/IBM/ITM
**********************************************************
Host  Prod  PID     Owner  Start  ID    ..Status
aix   ph    512006  itm1   Mar    None  ...running
aix   pk    430296                None
aix   px    237714                None
```

4.2.8 The SetPerm command

You can make use the `SetPerm` command to set file permissions to ensure that
the permissions are set to those permissions that appropriate for an applications
user. This command can only be run as the root user.

The correct method and usage of the SetPerm command was covered in
"Changing the file permissions" on page 92.

4.2.9 Best practices

The best practices for monitoring are:

► Planning

Have the monitoring requirement for the entire environment mapped out from
what agents will be deployed where and at what level the alert threshold
needs to be set, although this setting might require tuning over time, to the
basics of application file system creation. Plan the creation of applications
users and groups and the file permissions and ownership of the application
files.

What sort of diagnostic data do you collect in the environment, or do you
simply monitor systems, applications, the health, and the operations of
components within the environment? How is any ongoing maintenance to be
performed? Does the environment have a centralized point for the
management of systems and maintenance?

► Installation and Configuration

The AIX Enterprise Edition Launchpad has simplified the installation of IBM Tivoli Enterprise Monitoring in a System p environment as discussed in 3.3, "Installing and configuring IBM Tivoli Monitoring V6.2" on page 50. Following this method helps to greatly reduce the amount of time and complexity around actually getting the tool operational.

After the environment is being actively monitored, further customization of thresholds for alerts and event integration might be required.

4.3 Operating TADDM

This section describes how to operate a TADDM server. With a simple configuration of the scope and the access list, TADDM discovers a lot of components from target servers. It can provide the physical topology of the target servers as well as the software topology. However, customization is necessary if you want to extend the functionality, such as creating Business Applications and Business Services. There are additional features in TADDM, including versioning, showing change history, and reporting. We describe:

4.3.1, "Starting and stopping the TADDM server" on page 185

4.3.2, "Discover components" on page 186

4.3.3, "Customizing TADDM" on page 212

4.3.4, "Analytics and reporting" on page 228

4.3.5, "Performance considerations" on page 236

4.3.1 Starting and stopping the TADDM server

When you need to stop and start the TADDM server manually for any reason, go to the $TADDM_installed_directory/dist/bin directory and run **./control stop** and **./control start** as a TADDM admin user as shown in Example 4-4 on page 186.

Example 4-4   Start or stop TADDM server as a TADDM admin user

```
$ id
uid=205(cmdbadm) gid=3(sys) groups=1(staff)

$ cd /opt/IBM/cmdb/dist/bin
$ ./control stop
$ ./control start
```

4.3.2 Discover components

TADDM discovers and collects configuration information for the entire application infrastructure, identifying deployed software components, physical servers, network devices, virtual LAN, and host data used in a runtime environment.

Discovery is performed using sensors that are currently built and deployed as part of the TADDM product:

1. To get into the product console, bring up your browser and enter a URL with the host name of the TADDM server and port number 9430 as shown in Figure 4-12 on page 187, for example, `http://taddm10:9430`.
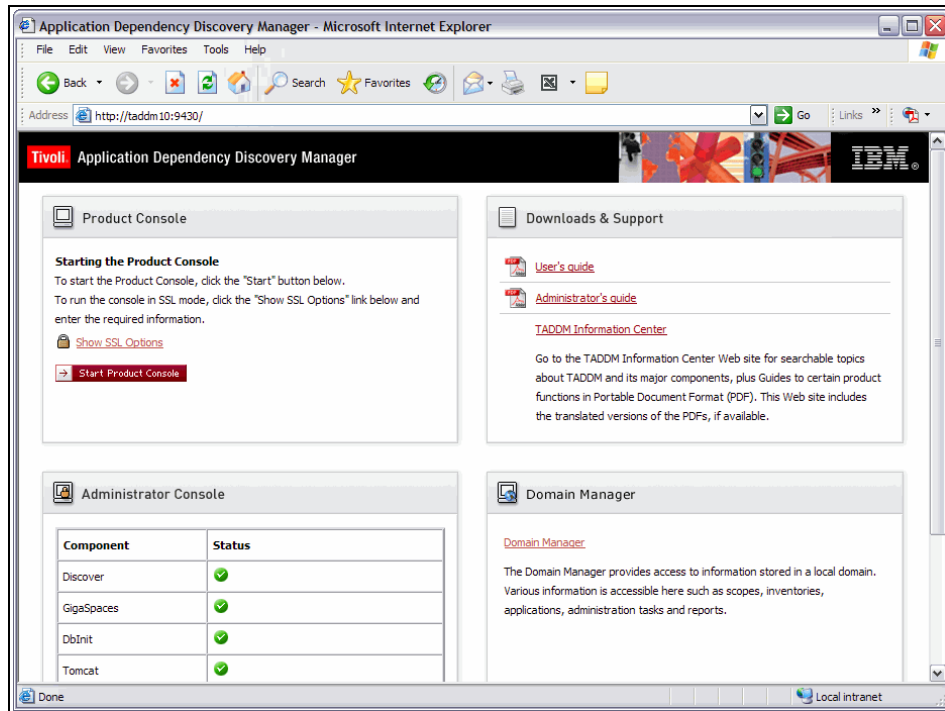
Figure 4-12   TADDM console

2. Clicking **Start Product Console** shows a window that asks you to open or save the configuration file as shown in Figure 4-13.
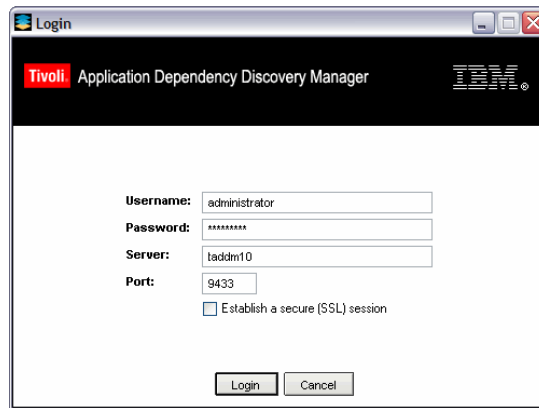
Figure 4-13   Open or save the configuration file

3. Select **Open** unless you want to save the file. If you save the file, it will be saved on your workstation for future use. If you click **Open** or click the previously saved file, it will show you the login panel as shown in Figure 4-14.

Figure 4-14   Login panel

The default user is `administrator` and the password is `collation`.

4. Enter the Username and Password (Figure 4-14). Click **Login** and the Tivoli Application Dependency Discovery Manager panel appears.

Scope

A *discovery scope* identifies the devices, computer systems, and other components in your infrastructure that you want the server to access. You specify scopes using IP addresses, ranges of IP addresses, or subnets to define the boundary of the networks that can be accessed during discovery. A scope can be as small as a single IP address or as large as a range of IP addresses or a Class C network. You can also exclude specific devices from the scope.

Before running discovery, you need to configure a discovery scope:

1. On the left pane, click the **Discovery** tab → **Scope**. The Scope pane is displayed.

2. To define a new discovery scope set, click **Add Set**. The Scope Set Name window is displayed as shown in Figure 4-15.
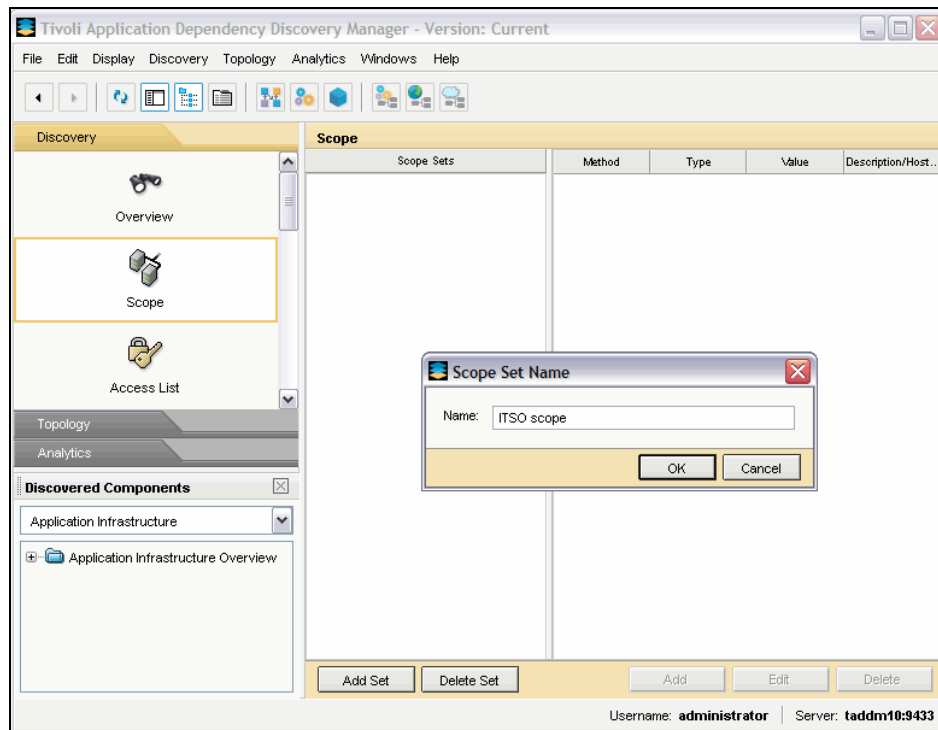
Figure 4-15   Scope Set Name

3. In the Name field, type the name for the new scope set.

4. Click **OK**. The new scope set appears in the Scope Sets list.

5. To add contents to the scope set, select the scope set that you have just created and click **Add**. The Add Scope window is displayed as shown in Figure 4-16.

Figure 4-16   Adding scope

To add the settings for the scope, complete one of the following steps:

– Select Subnet from the IP Type list and type the IP address of the subnet in the IP Address field.

– Select Range from the IP Type list and type the starting IP address and the ending IP address in the IP Addresses field.

– Select Host from the IP Type list and type the IP address of the host in the IP Address field or type the host name in the Hostname field.

To exclude devices from your scope, click Add Exclusion and complete one of the following steps:

– From the IP Type list, select Subnet and type the IP address of the subnet in the IP Address field.

– From the IP Type list, select Range and type the starting IP address and the ending IP address in the IP Address field.

To save the scope, click **OK**. The new scope appears in the list.

Access List

The *Access List* is a collection of all user names, passwords, and Simple Network Management Protocol (SNMP) community strings that the server uses when accessing entities, which are stored as configuration items in the database, in your infrastructure. You have to set up this list for the configuration items that you want to discover.

To add a new Access List entry, complete the following steps from the product console:

1. On the left, click the **Discovery** tab → **Access List**. The Access List pane is displayed.

2. To add a new computer system into the Access List, click **Add**. The Access Details panel is displayed as shown in Figure 4-17.
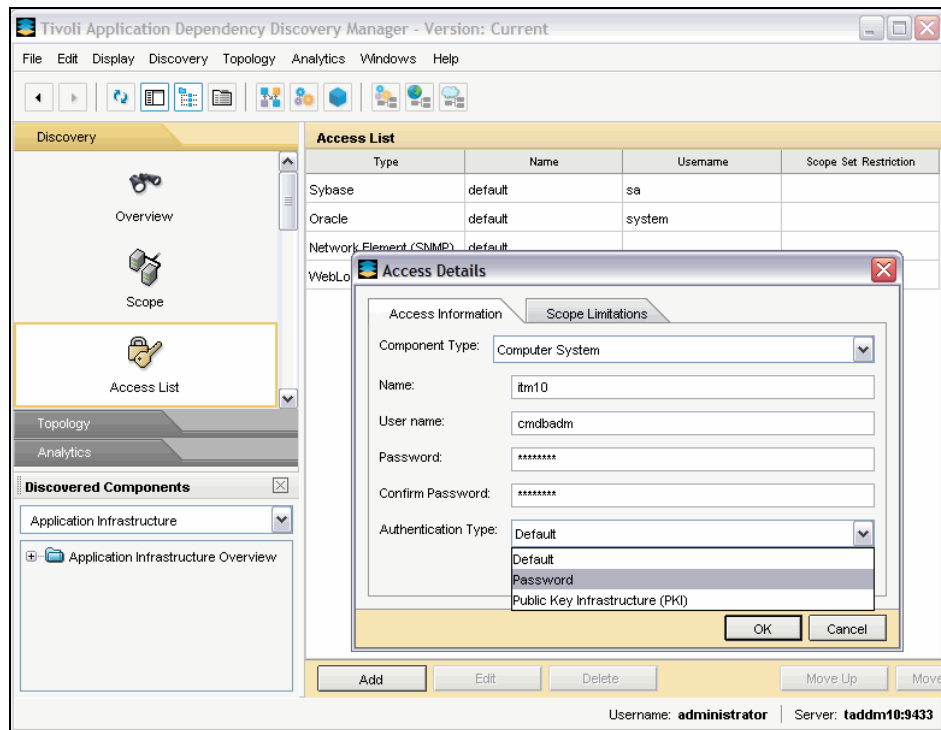
Figure 4-17   Adding a computer system

3. From the Component Type list, select the component type that you want to discover.

4. For all component types other than Network Element (SNMP), complete the following steps:

   – In the Name field, type the name of the Access List entry.

   – In the User name field, type the user name to log in to the component that you want to discover. Notice that we already added `cmdbadm` user in all the target servers for discovery before adding computer systems here.

   – In the Password field, type the password to log in to the component that you want to discover.

   – In the Confirm Password field, retype the password for confirmation.

   Additional steps might be required based on the component type that you have selected. We selected `password` as an authentication type for computer systems type.

   Note: TADDM attempts to use the credentials in the order that they appear in the access list. To speed up discovery, make sure that you have the credentials in the order that makes the most sense for your environment. Order the credentials from the most general to the most specific. Use the list order in combination with limiting scope. If a credential only applies to a particular scope set, be sure to limit its use to that scope set.

5. In the next window, we added a database in the Access List as shown in Figure 4-18.
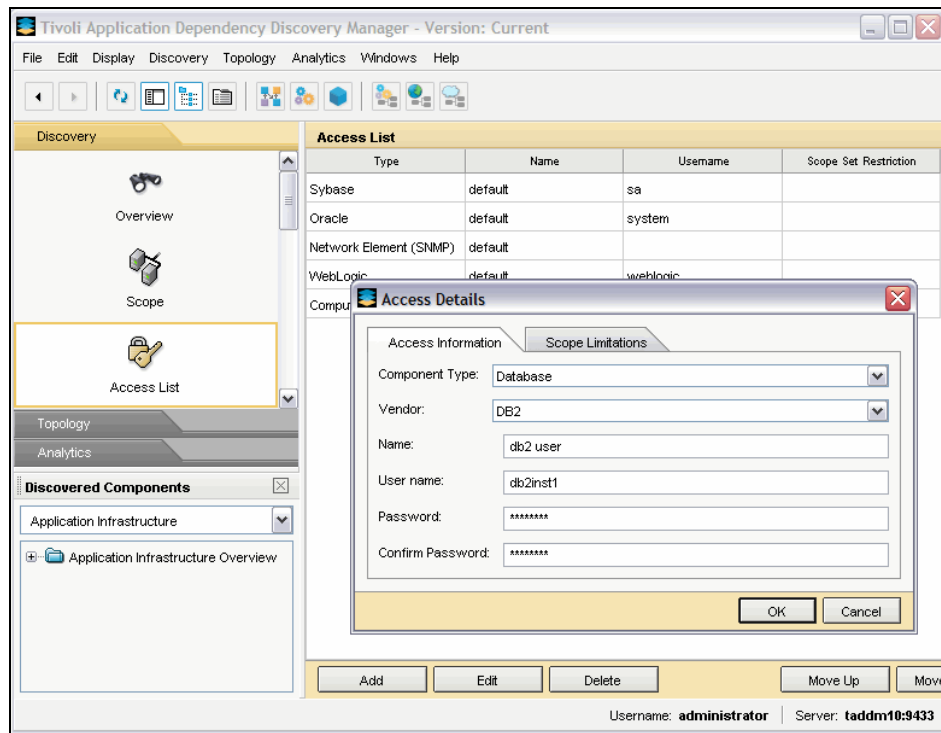
Figure 4-18   Adding a database

6. We took the same steps to add two HMCs and four more servers. Figure 4-19 shows the list of computer systems and the database that we have created.
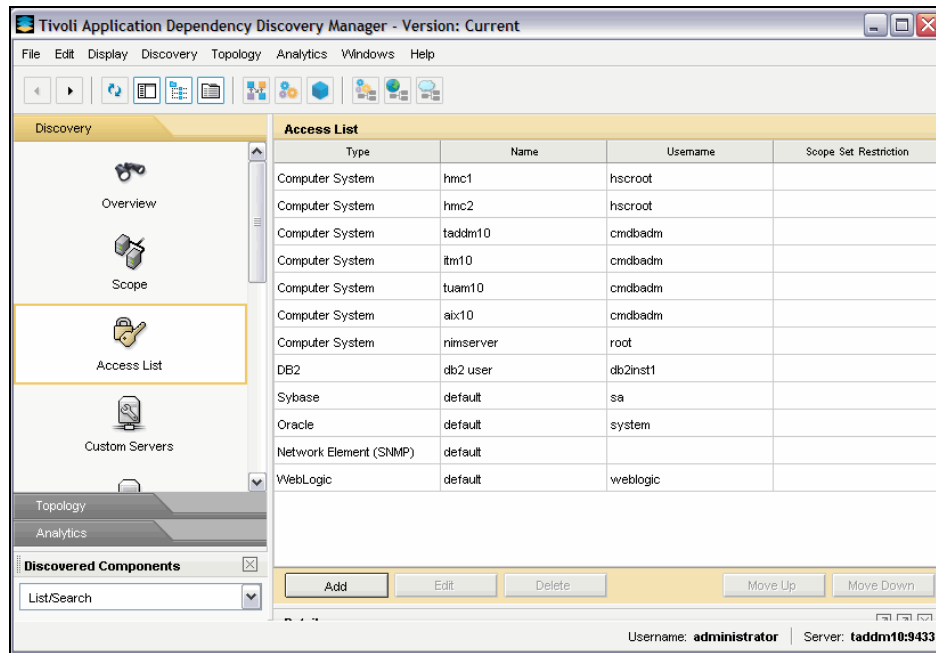
Figure 4-19   The list of computer systems and the database

Discovery profile

A *discovery profile* is a set of rules that controls how discoveries run. Using discovery profiles, you can control what TADDM discovers. For example, you can configure individual discovery sensors, manage various configurations for the same sensor, select the appropriate sensor based on a set of criteria, and control sets of different sensor configurations to be applied on a single discovery run.

By default, there are three levels of discovery profiles, which are described in Table 4-2 on page 195.

Table 4-2   Discovery profiles by default in TADDM

Level	Type of discovery	What is discovered
Level 1	Credential-less	Basic information about active computer systems: host name, fully qualified domain name (FQDN), OS release level, IP address, and open ports. Network operating systems are also discovered.
Level 2	OS credentials only	Detailed information about all of the operating systems.
Level 3	Full credential	Entire application infrastructure: deployed software components, physical servers, network devices, virtual LANs, and hosts.

In order to run discovery:

1. Click the **Discovery** tab → **Overview**, and click **Run Discovery**.

2. Check **ITSO scope** under Scopes, and select **Level 3 Discovery** from the Profile list box. Click **OK** as shown in Figure 4-20 on page 196.
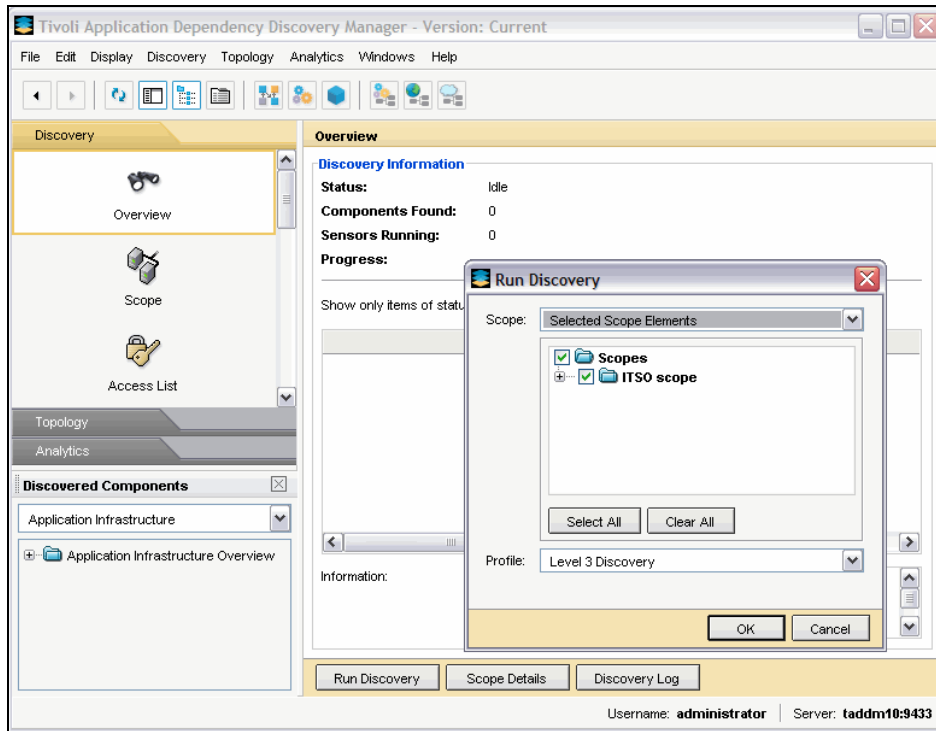
Figure 4-20   Running discovery

3. While running discovery, click **Discovery Log** to see Discovery Log Details as shown in Figure 4-21.
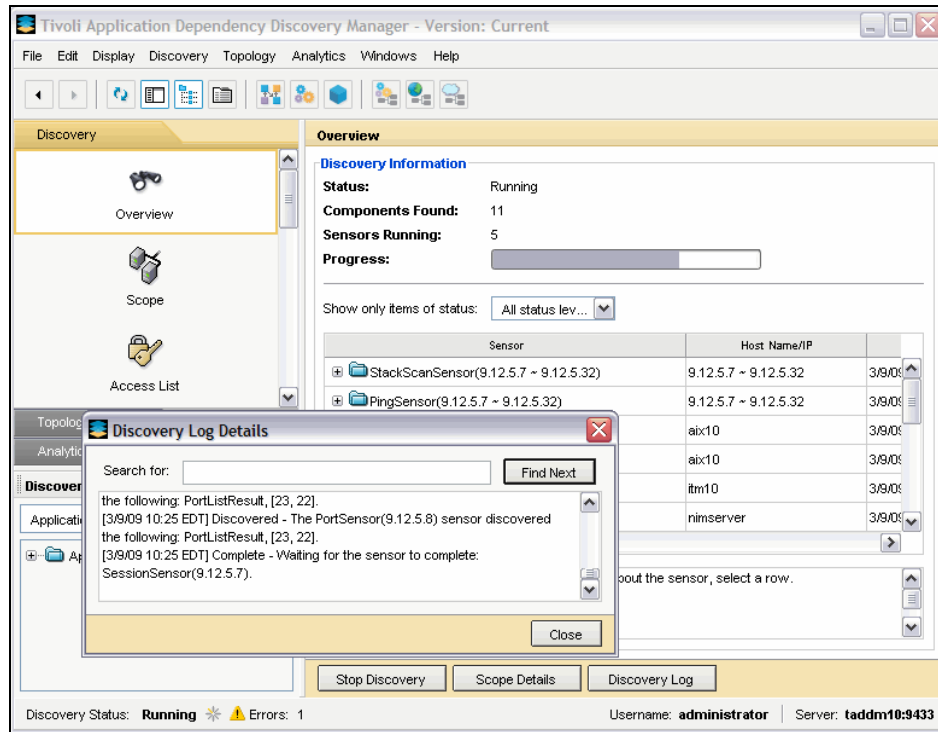
Figure 4-21   Discovery Log Details

When the discovery finishes, a line of information appears in red on the bottom of the window (Figure 4-22 on page 198): **"`Changes have occurred. Reload the view.`"**

4. Click **Display** → **Reload** in the product console to apply the changes as shown in Figure 4-22.
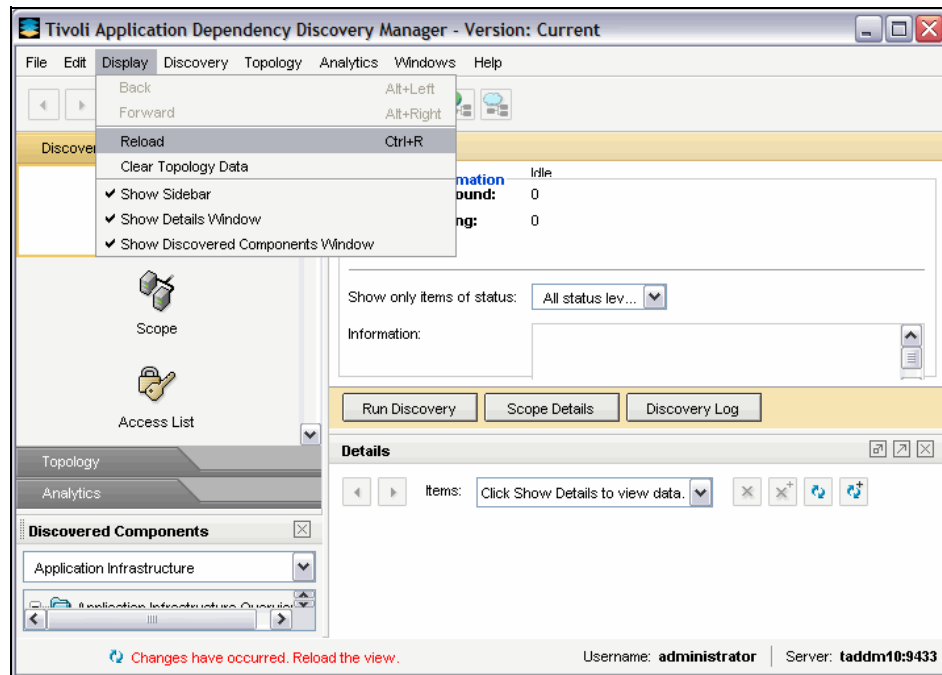
Figure 4-22   Reload the view

Now, you are ready to see the components that TADDM has just discovered or the *topology* of the components. In order to search the components, click **Discovered Components** in the left pane as shown in Figure 4-23.
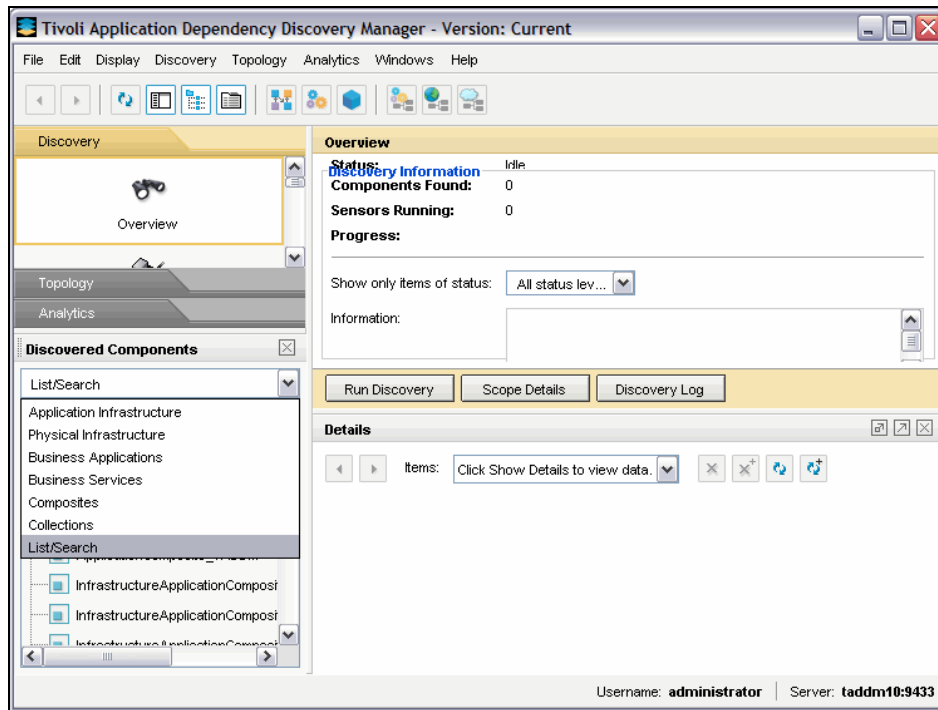
Figure 4-23   List search

5. Click **List/Search** and move down to select **component.**

6. Select **IBM AIX Computer System** to show the list of servers as shown in Figure 4-24.
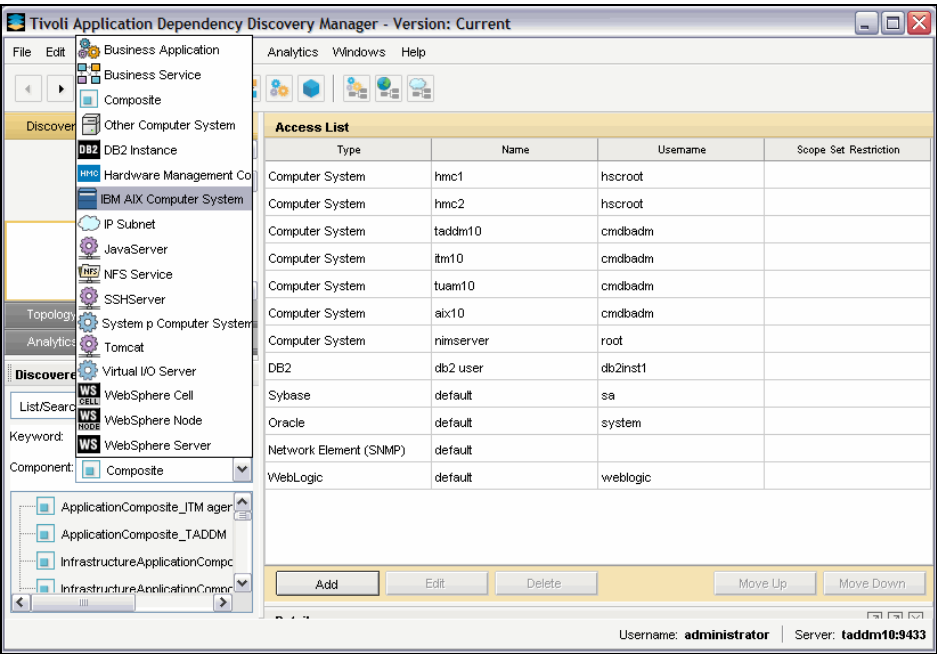
Figure 4-24   Selecting IBM AIX Computer System

7. We have five logical partitions on AIX as shown in Figure 4-25. Note that if you have a logical partition and you do not see the `LPAR` symbol to the left of the component name, verify with the HMC that the LPAR is connected and has been added properly.
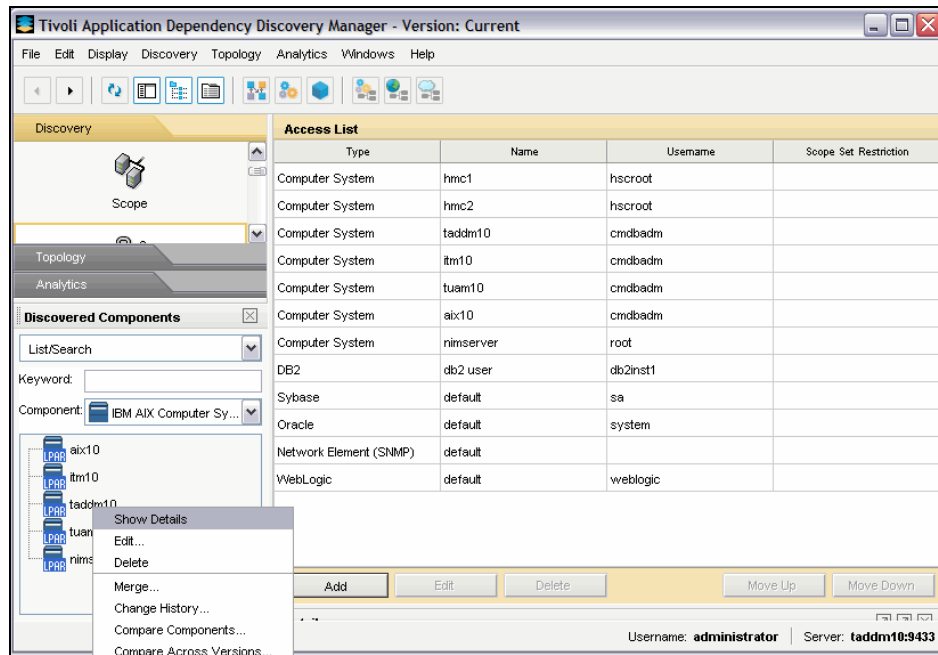
Figure 4-25   Show Details

8. Right-click one of the servers, and then click **Show Details**.

    You see the details of the server as shown in Figure 4-26. The details include general information about the hardware, OS, storage, file systems, IP, installed filesets, patch level, and the contents of major configuration files.
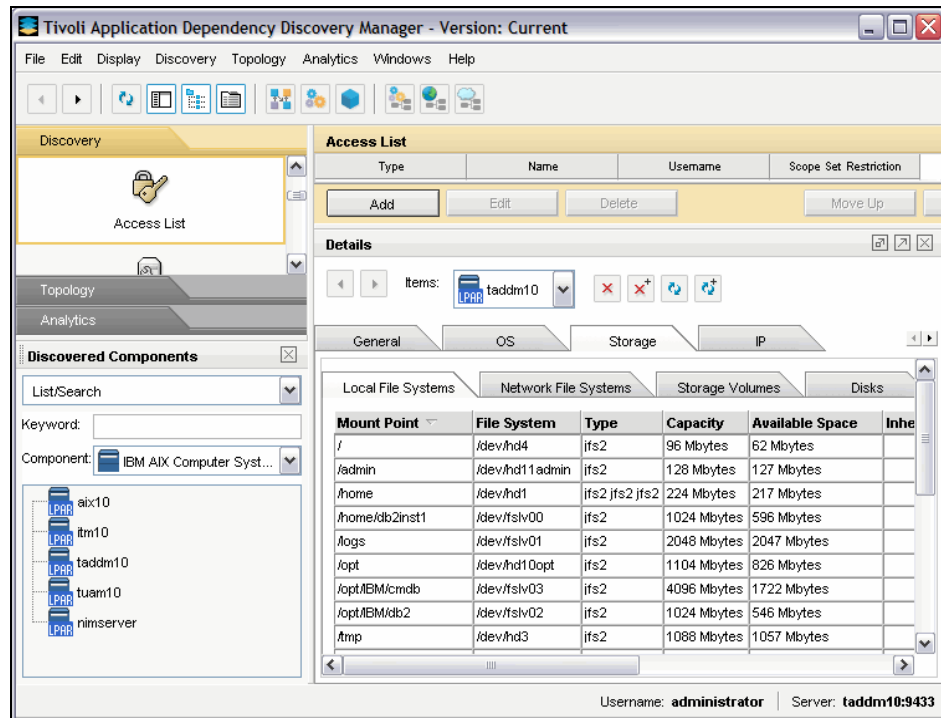
Figure 4-26   Details of a computer system

9. In the same manner, we selected **DB2 Instance** to see the details of a database as shown in Figure 4-27.
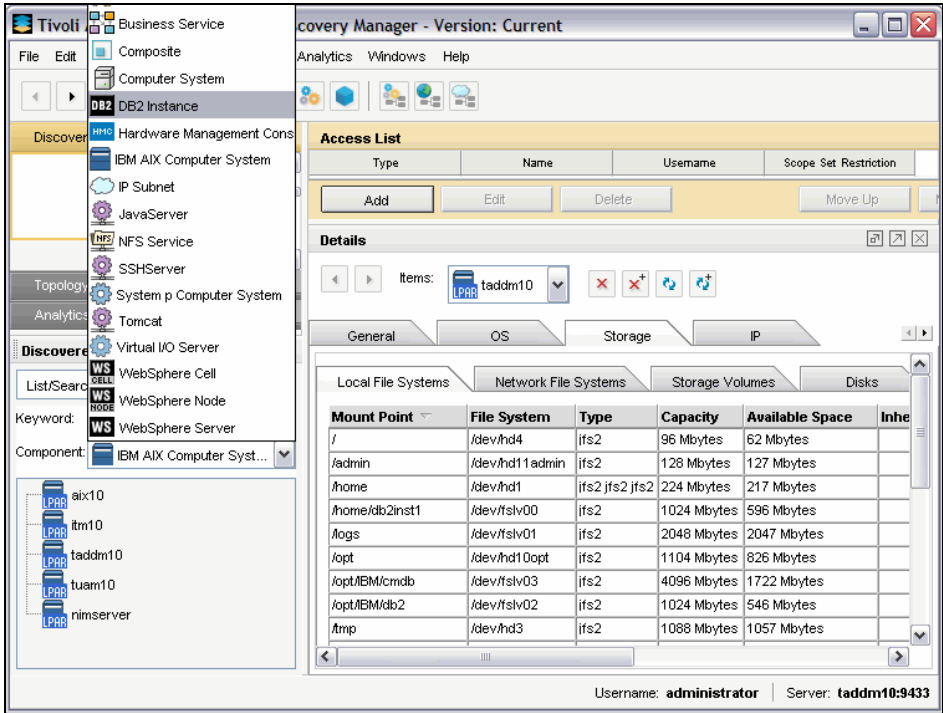
Figure 4-27   Selecting a DB2 instance

10. Right-click one of the DB2 instances and click **Show Details** as shown in Figure 4-28. The details include database information as well as the database instance itself.
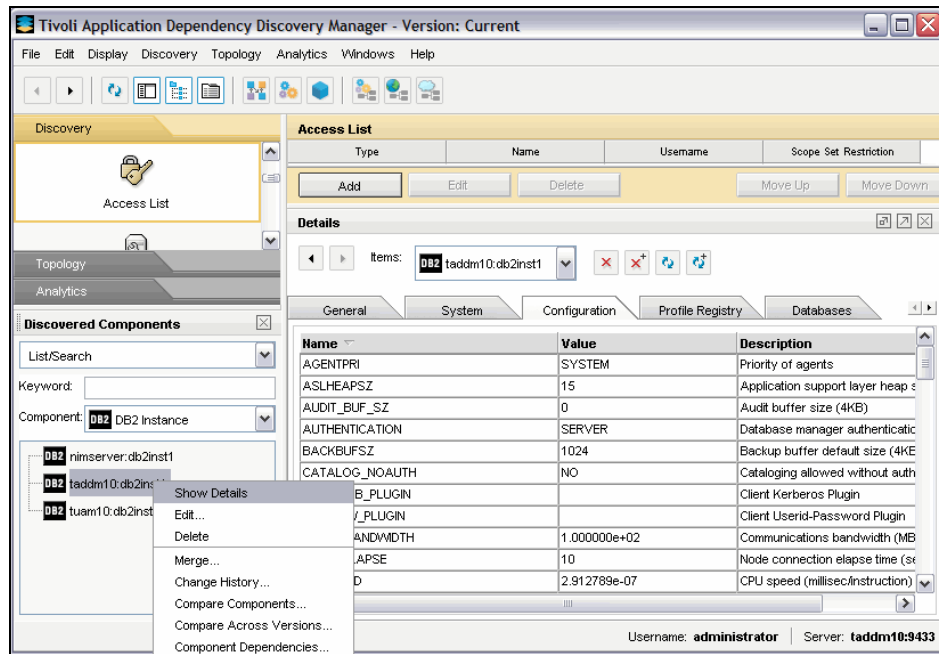
Figure 4-28   Details of a database instance

11. Now, click the **Topology** tab in the left pane to check the relationship between components. Click **Application Infrastructure** to see how the applications are related as shown in Figure 4-29.
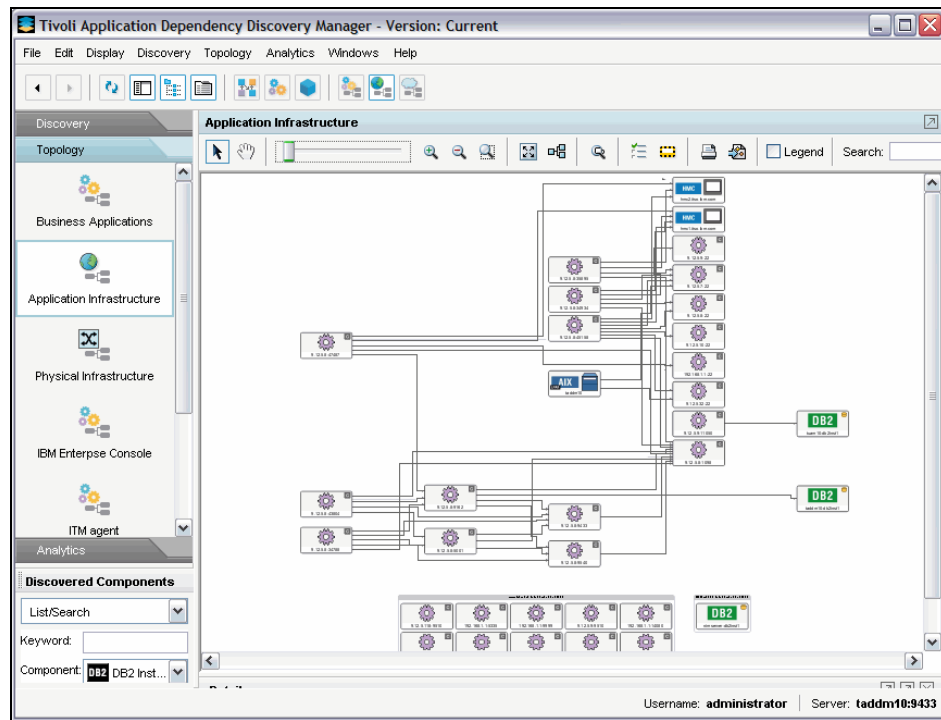
Figure 4-29   Topology for Application Infrastructure

You can zoom in or out on the Topology window Application Infrastructure by sliding the zoom indicator (see the cursor in Figure 4-30). Or, just click the zoom slider and use the arrow keys for sizing the window.
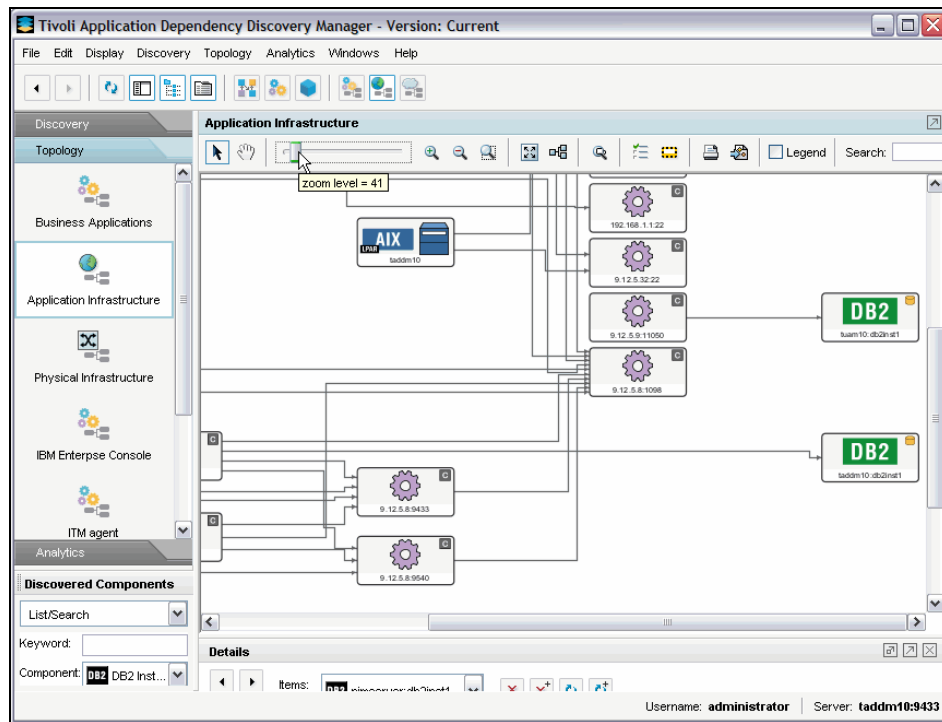
Figure 4-30   Zooming in on topology

Another view provided in the Topology tab is the Physical Infrastructure. Click **Physical Infrastructure** to see how the components are physically associated.

12.Choose a network and right-click **Show Details** as shown in Figure 4-31.

Even though the IP addresses of our servers are in the network of 9.12.5.x, they are shown in the network of 9.12.4.0, because our subnet mask is 255.255.252.0.
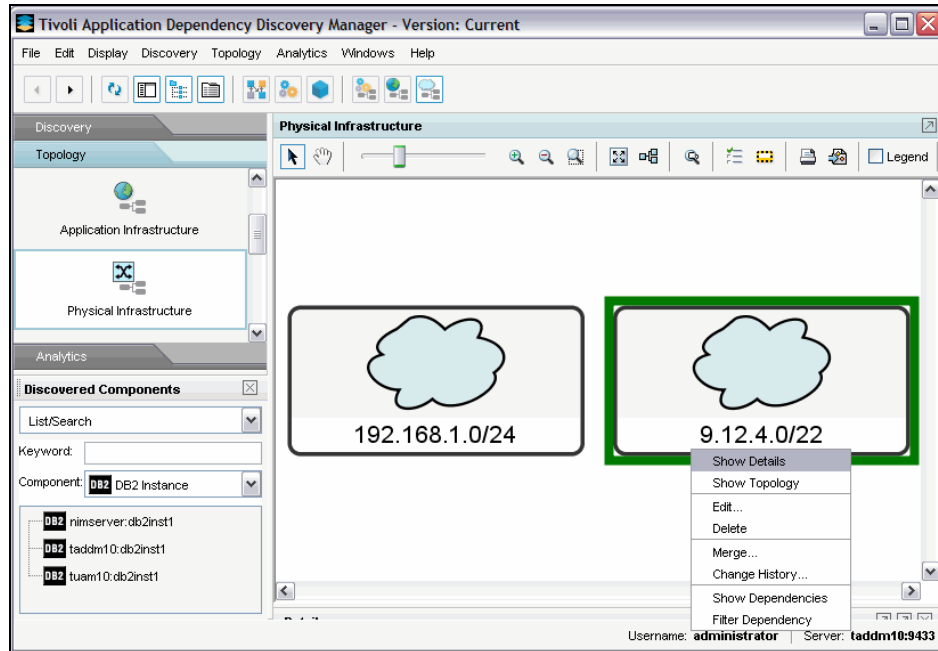
Figure 4-31   Physical Infrastructure

13. The servers and HMCs in the network are displayed. Right-click one of the servers to show details as shown in Figure 4-32.
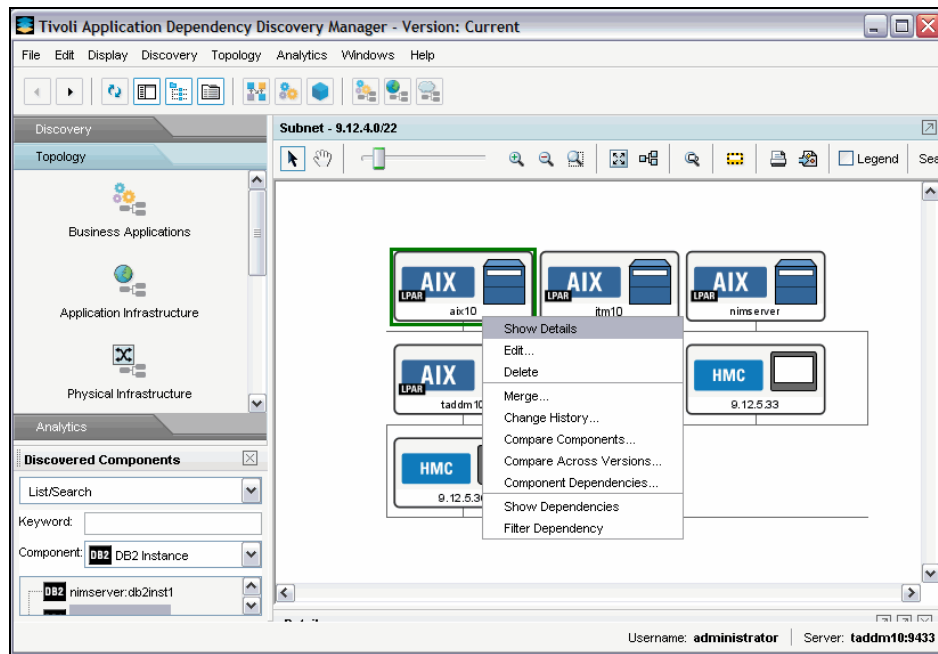
Figure 4-32   Showing the details of a server

Schedule discovery

You can initiate TADDM discoveries on demand based on a schedule. Scheduling is a methodical way to keep the TADDM database up-to-date. You can schedule discoveries to start during periods of low activity, such as on weekends, during times when batch jobs are likely to run, such as late at night, or during periods of high activity, such as mid-morning or mid-afternoon. You can schedule discoveries before and after normal maintenance windows, enabling you to quickly check what changes were made during that interval.

To create a new discovery schedule, complete the following steps:

1. On the Discovery tab, click **Schedule**.

2. Select **Add** to define a new discovery schedule. The Discovery Schedule window opens as shown in Figure 4-33. The date and time that are shown in the window reflect the date and time on the TADDM server.
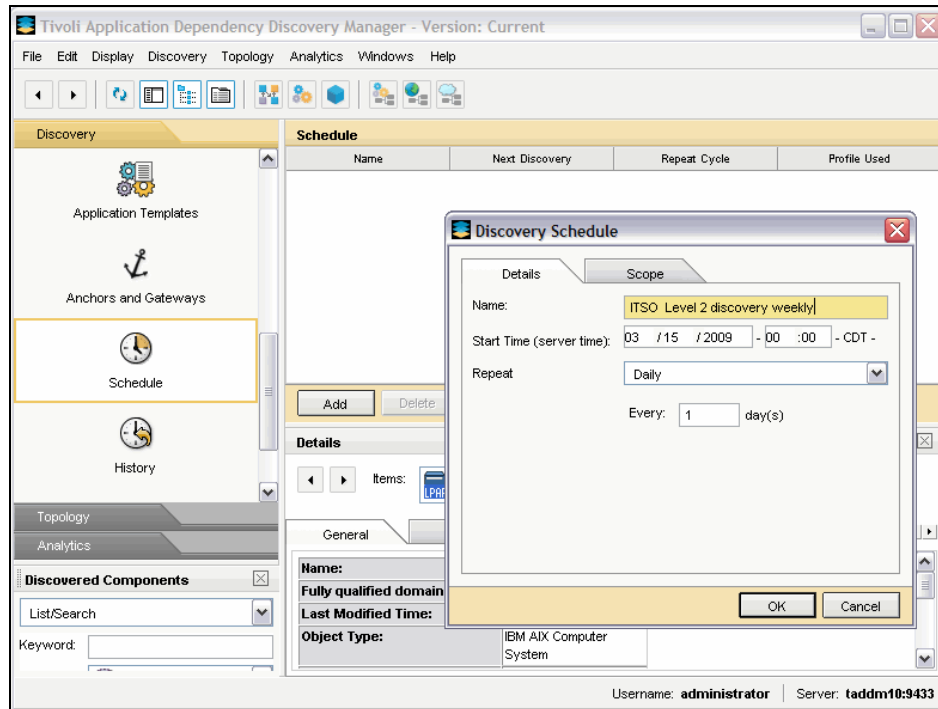
Figure 4-33   Schedule discovery

3. Provide a name, starting time, and repeat interval. And then, click **Scope** to specify a scope and the level of discovery as shown in Figure 4-34.
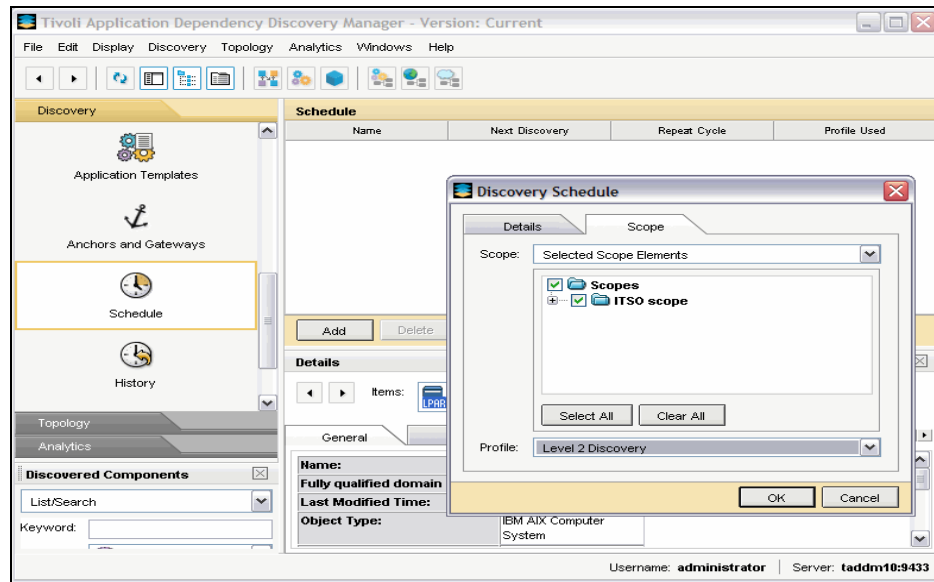
Figure 4-34   Defining the scope in the discovery schedule

History of discovery

If you want to see the previous discoveries, click **History** in the Discovery tab shown in Figure 4-35.
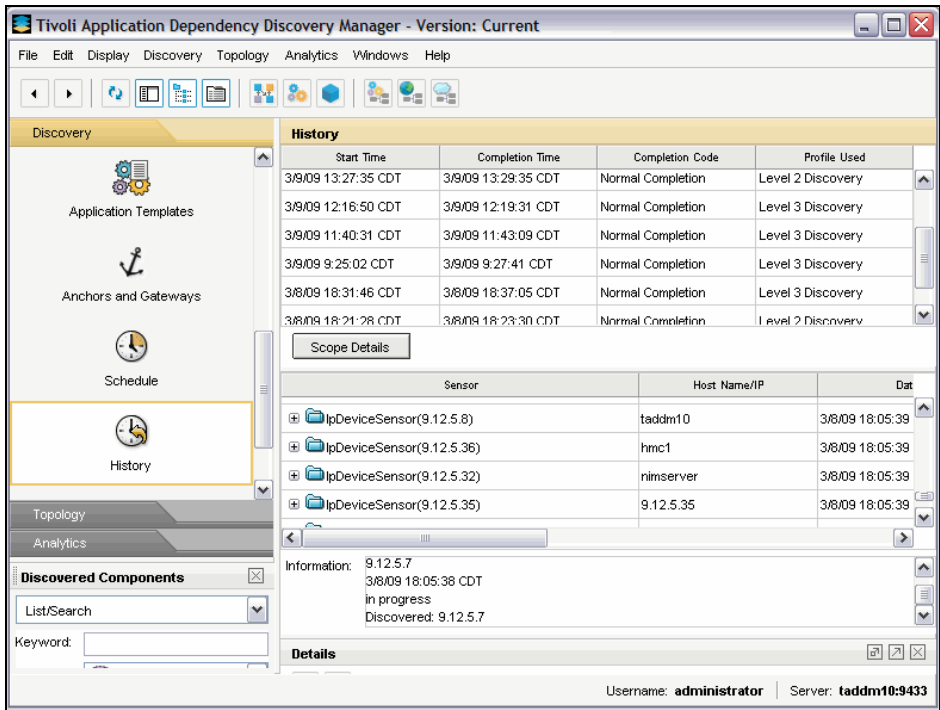
Figure 4-35   History of discovery

Versions

Versions are snapshots of the current infrastructure. Versions are read-only views of the entire topology. Analytic reports support comparisons between versions.

To create a new version, click **Versions** in the Discovery tab, and select **Create**. Enter a version name, and click **OK** as shown in Figure 4-36 on page 212.

> **Note:** We suggest that you have naming rules for version names to avoid possible confusion later. For example, you can use the date as part of the version name. A version ID is generated in TADDM whenever a new version is created. The current version has an ID value of 0.
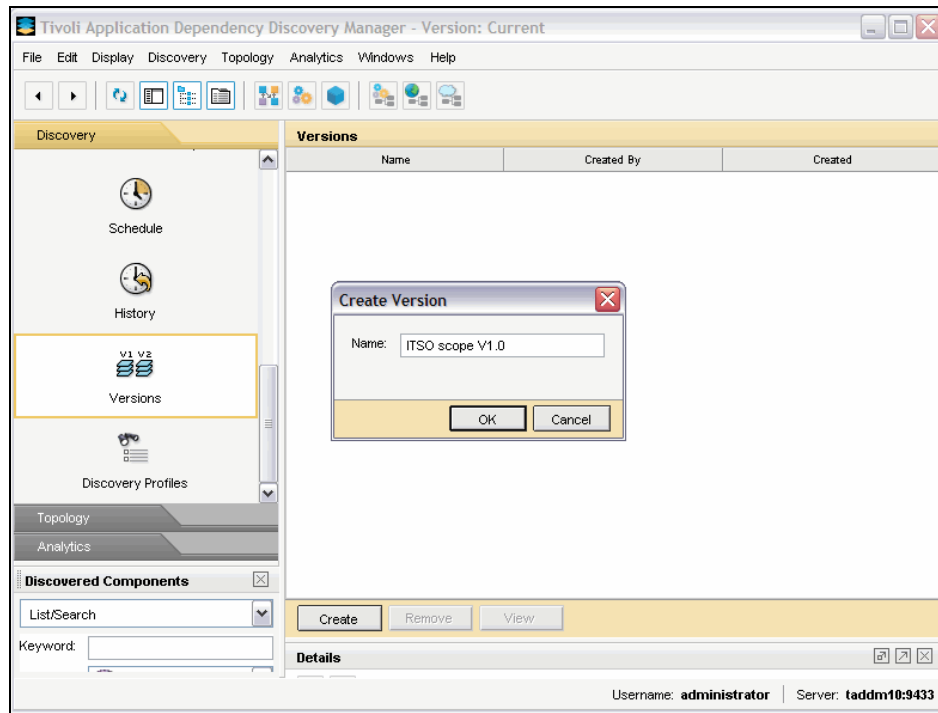
Figure 4-36 Creating a version

4.3.3 Customizing TADDM

In this section, we describe how to add a custom server and a Business Application.

Custom server

You can create custom servers to discover and categorize servers that are not, by default, supported by TADDM. Your infrastructure might contain software applications and server types, such as custom Java servers, that are not automatically categorized by TADDM. Any server process with a TCP listening port that is not recognized is categorized into an Unknown Server category. Unknown servers are not displayed in the Topology tab and cannot take advantage of most of the functions. You do, however, receive basic information, such as the name and the runtime data of the unknown server.

You can define a custom server to create a template that sets up the membership rules for the custom server. During a discovery, any unknown server is automatically categorized as a custom server of this type if the runtime

information matches the criteria that you have defined in the template. Custom servers are displayed in the topology, and you can view details about them. Although these details are not as complete as the details that are provided for supported servers, defining custom servers allows all of the components in your infrastructure to participate in the topology view and comparisons. You can manage custom servers by selecting the **Discovery** tab → **Custom Servers**:

1. Before adding a server, run a basic discovery to check for unknown servers. You can run a report on unknown servers to help you identify patterns to use in the custom server template.

2. To identify patterns in unknown servers, click the **Analytics** tab → **Inventory**, and the Inventory pane is displayed in the workspace. Then, select **Unknown Servers**, and click **Run Report** as shown in Figure 4-37.
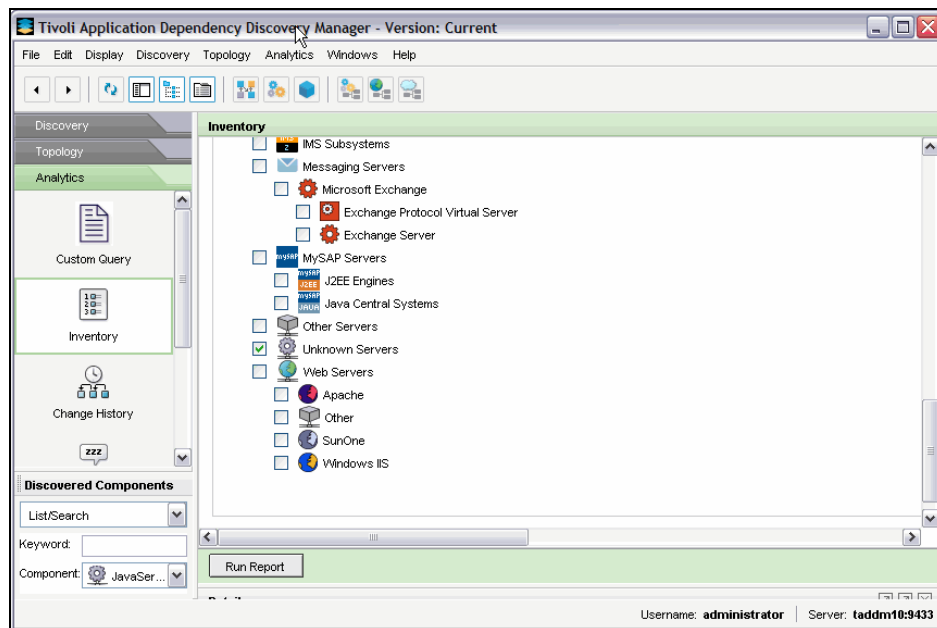
Figure 4-37   Collecting unknown servers

You can identify a pattern in the configuration of the unknown server, such as the program name, arguments, environment, and port. Use this pattern to create the identifying criteria for the custom server template.

3. In our environment, we use sendmail to create a new custom server. We highlight **sendmail** and click **Create Custom Server** as shown in Figure 4-38.
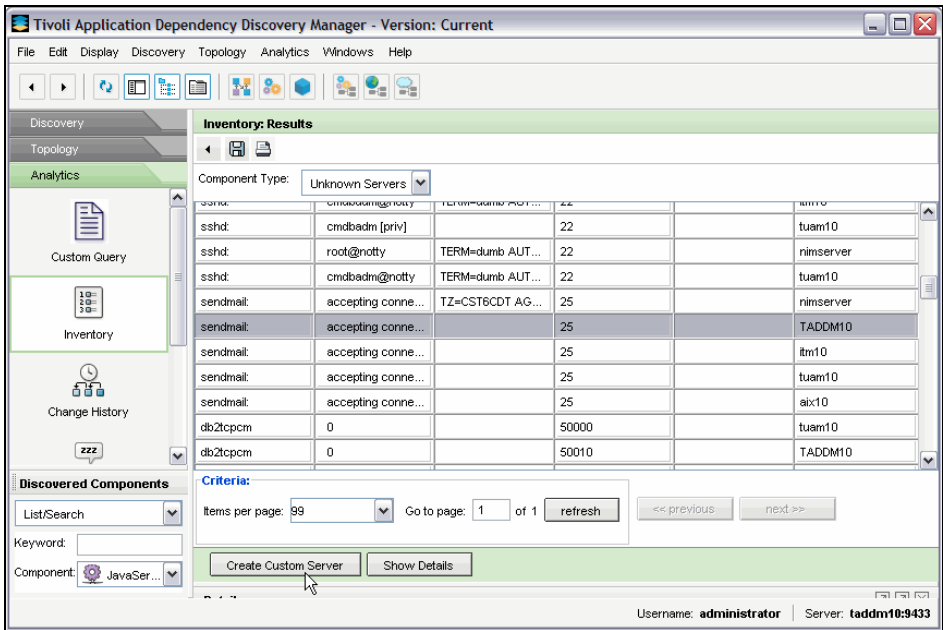
Figure 4-38   Identifying unknown server patterns

The Custom Server Details panel General Info & Criteria tab is displayed (Figure 4-39). In the Name field, type the name of the custom server. In this example, we used `Sendmail - email server`.

4. To enable the custom server definition, select **Enabled**. And then, we changed the icon for sendmail. Define the criteria for the custom server. If you are creating a custom server from an unknown server, you will see at least one criterion that is set up. Figure 4-39 shows our example.
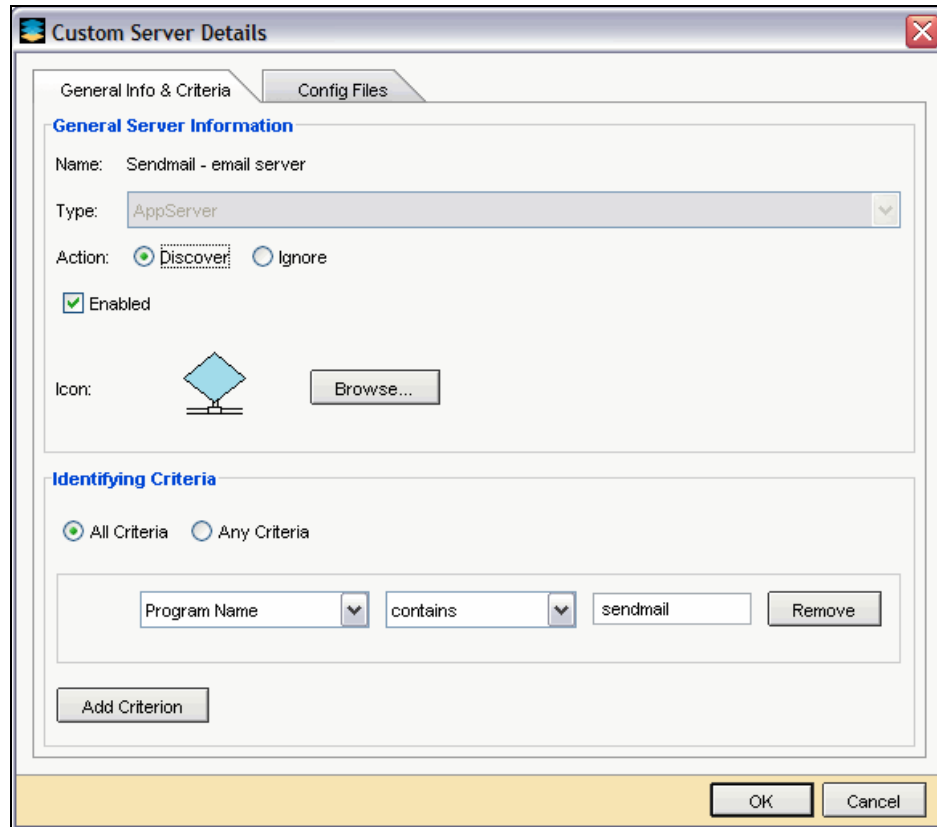
Figure 4-39   Creating a custom server

5. To add configuration files, click the **Config Files** tab. The Config Files page is displayed. On the Config Files page, click **Add**. The Search Path for Capture File window (Figure 4-40 on page 216) is displayed. From the Search Path list, provide a directory path for the configuration file.

To capture the contents of the configuration file, click **Capture file contents**, and optionally, specify the maximum number of bytes of the captured configuration file. To recursively explore through the directory structure to search for the specified file, click **Recurse Directory Content**.

To save the settings for your custom server, click **OK**. Refer to Figure 4-40.
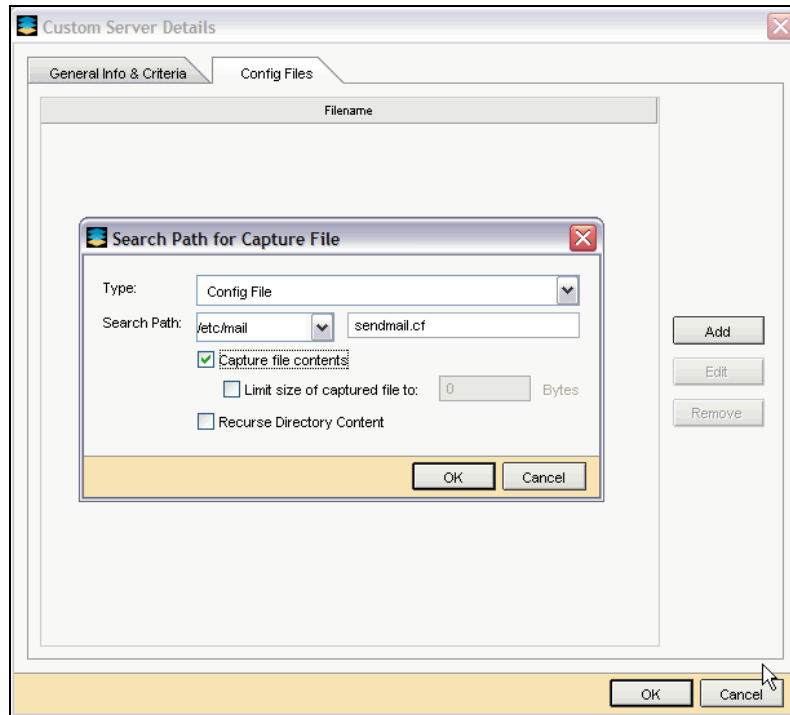
Figure 4-40   Selecting config files

6. To reposition entries in the Custom Servers pane, select **sendmail** and move it up to your desired position as shown in Figure 4-41. Note that you can also choose to enable or disable discovery for specific applications in the Custom Server pane. Click **Save** to save the changes to be applied.
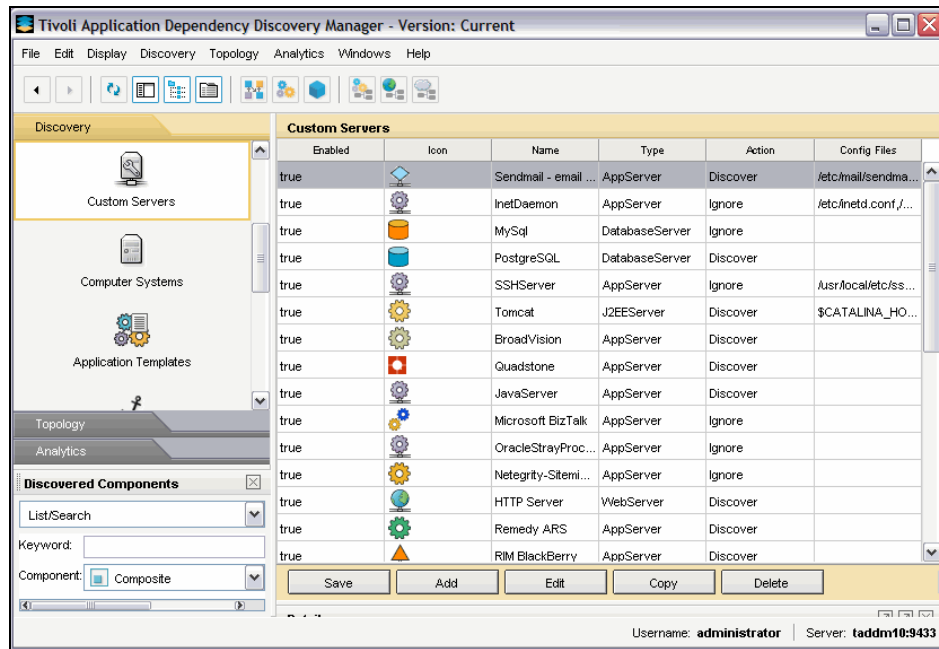
Figure 4-41   Repositioning a custom server

Now, we can run a new discovery and check our new custom server that we have just created. Run discovery and check the results by clicking the **Topology** tab → **Application Infrastructure** as shown in Figure 4-42 on page 218. The applications with port number 25 are sendmails. The sendmail applications are marked with the graphic that is shown here to the left.
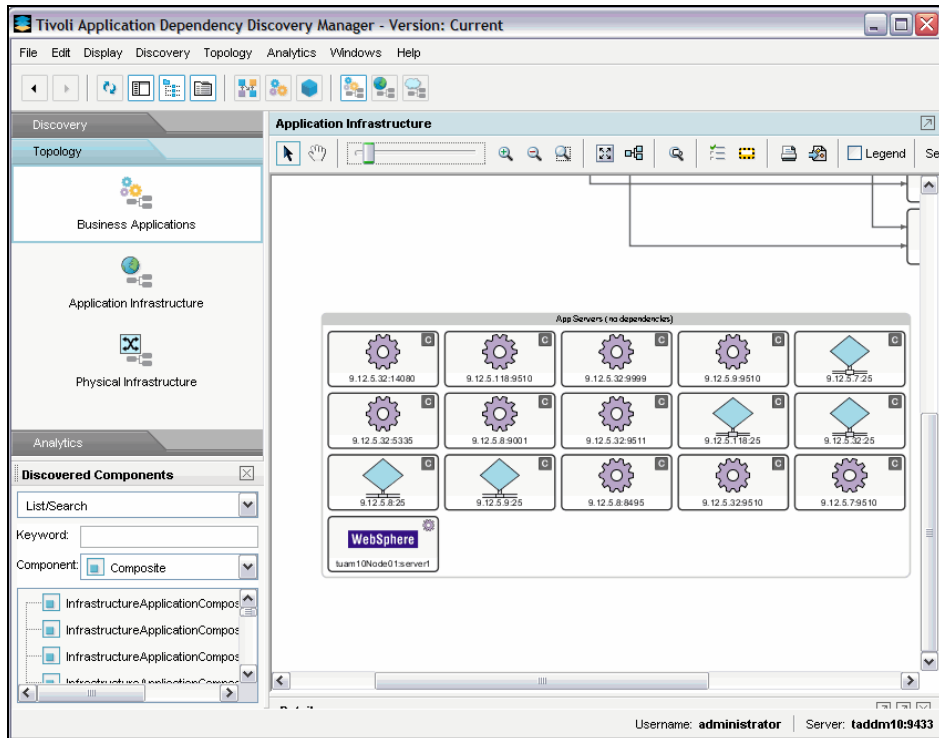
Figure 4-42   Sendmail custom server discovered

If you want to see more details, right-click the discovered component and click **Show Details**. You can see all of the information related to the component, including the configuration file.

Business Applications and Business Services

A *Business Application* is the way to group the various kinds of IT resources into a logical group, and this logical group acts together as one unit to provide a type of service. You can treat Business Applications as units for discovery and analytics.

The top level in the component hierarchy of TADDM is the *Business Service*. Business Services can contain any number of the lower level resources, from Business Applications to modules in a WebSphere server or specific configuration files on systems. The purpose of the Business Service is to consolidate multiple lower level objects and their relationships in order to perform reporting and analysis, considering all related resources.

1. To create Business Applications, click **Edit** → **Create Business Application**. Enter a name (we entered `ITM agent`) and description (we entered `IBM Tivoli Monitoring`) as shown in Figure 4-43.
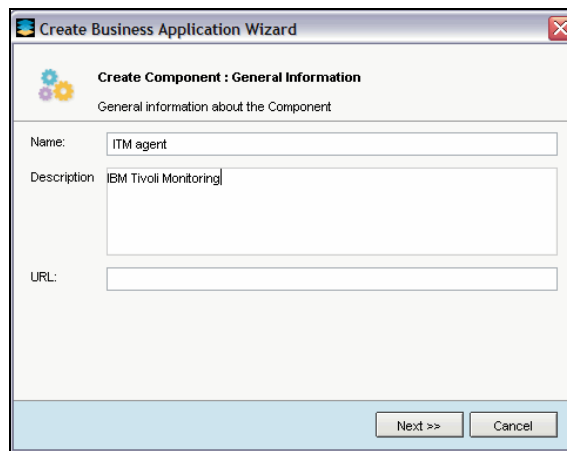
Figure 4-43   Creating a Business Application

2. Click **Next** to select the components. We selected a DB2 instance and one Java server to show you an example of a Business Application as shown in Figure 4-44. You can choose applications that are appropriate for your environment.

   When you manually add a component to a Business Application, you need to choose a functional group. The *functional group* defines what components are compared when a comparison is performed between two Business Applications. We selected the default functional group in this example.
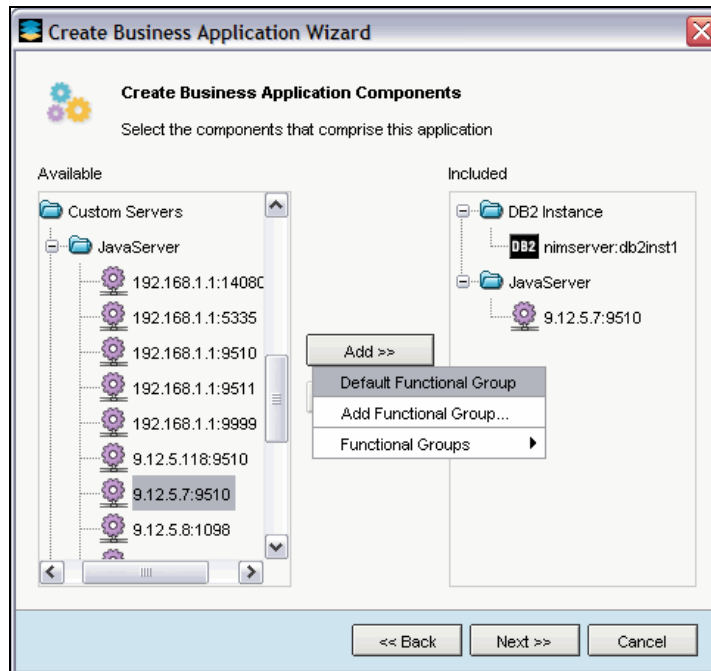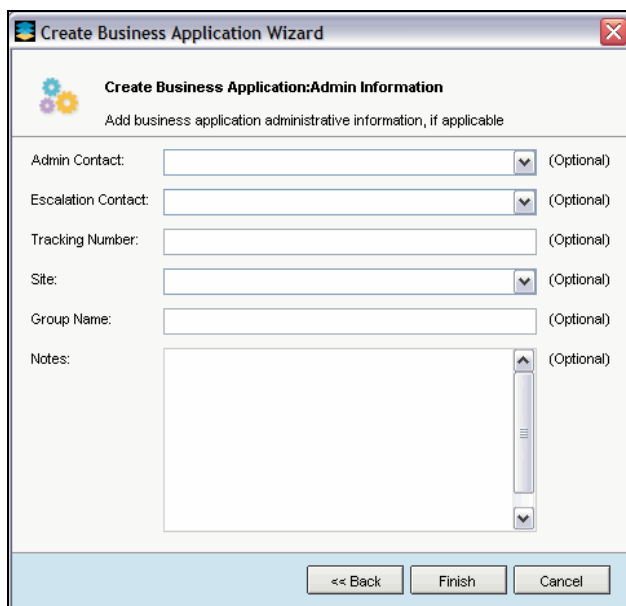
Figure 4-44   Assigning components for the Business Application
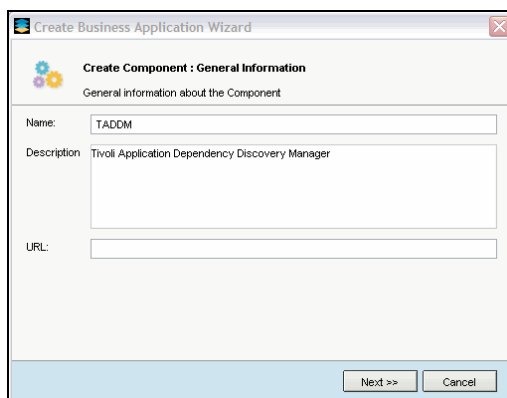
In the next window, which is shown in Figure 4-45, you can provide administrative information. We skipped this page, because it is optional.

Figure 4-45   Administrative information

3. In the same manner, we created the second Business Application, which is named `TADDM`, as shown in Figure 4-46.

Figure 4-46   Creating a second Business Application

4. For the second Business Application, we assigned one DB2 instance and three Java servers to this application as shown in Figure 4-47.
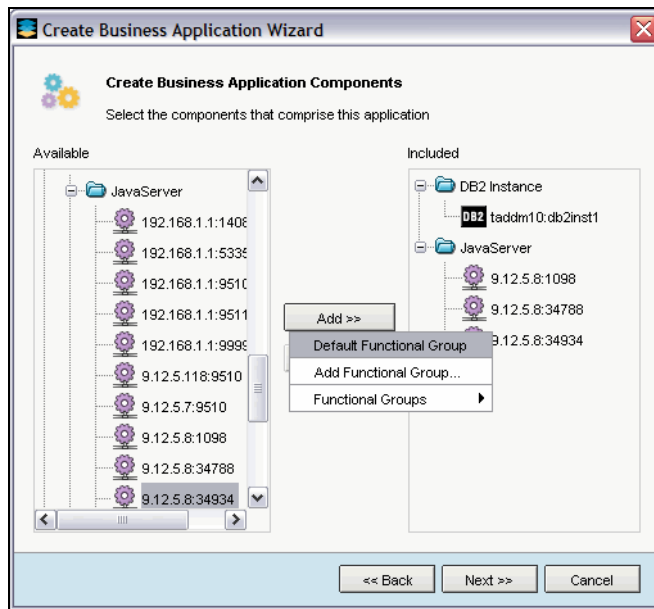
Figure 4-47   Assigning components for the second Business Application

5. After creating two Business Applications, go to the **Topology** tab and click **Business Applications**. You will see the two Business Applications that you have just created as shown in Figure 4-48.
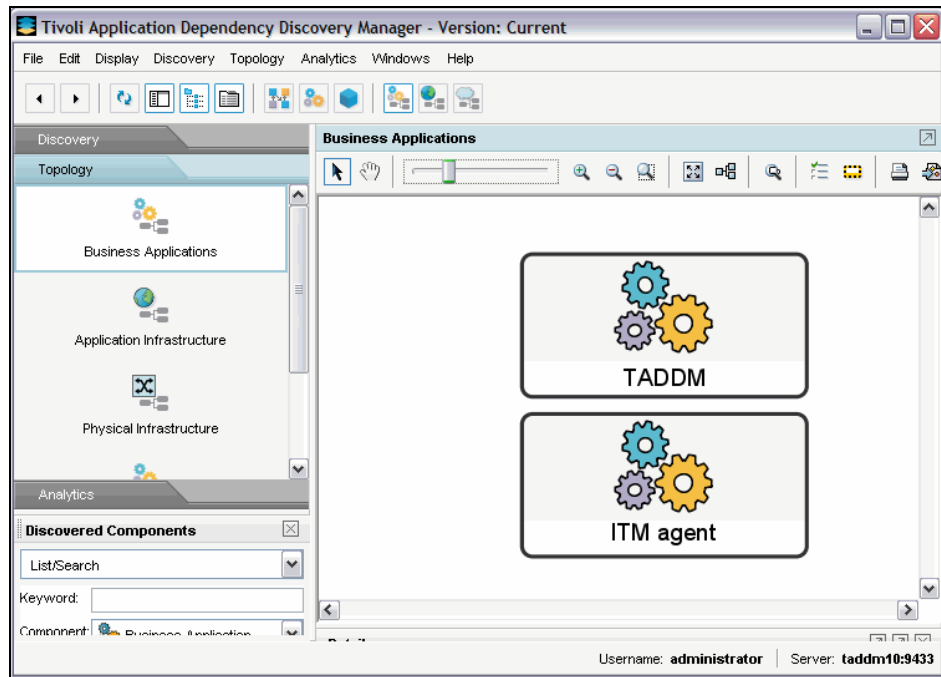
Figure 4-48   Showing Business Applications

Now, we create a Business Service that includes the two Business Applications that we have just created.

6. Click **Edit → Create Business Service** as shown in Figure 4-49.
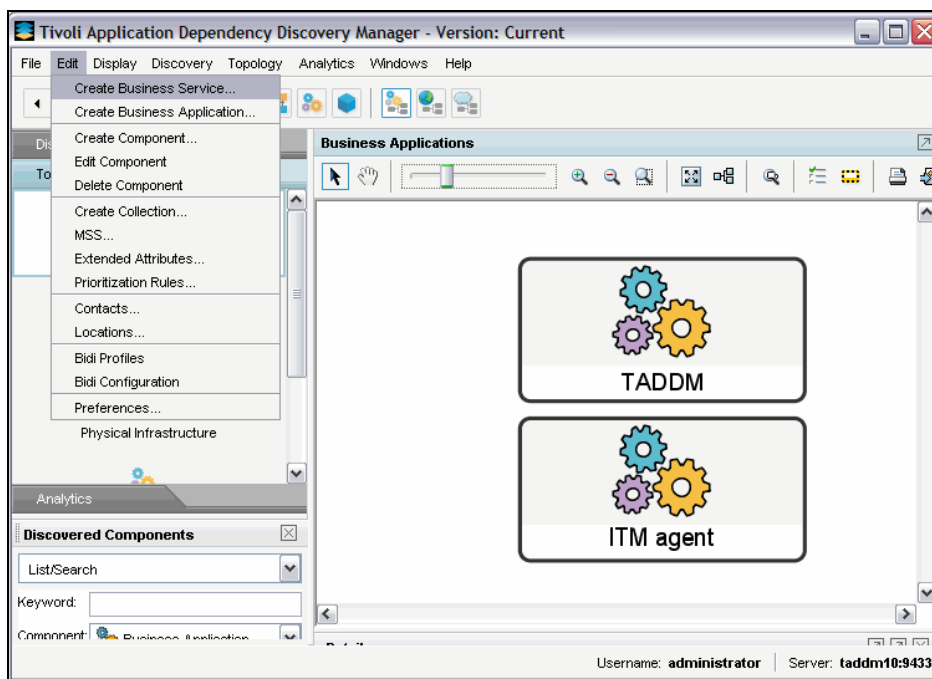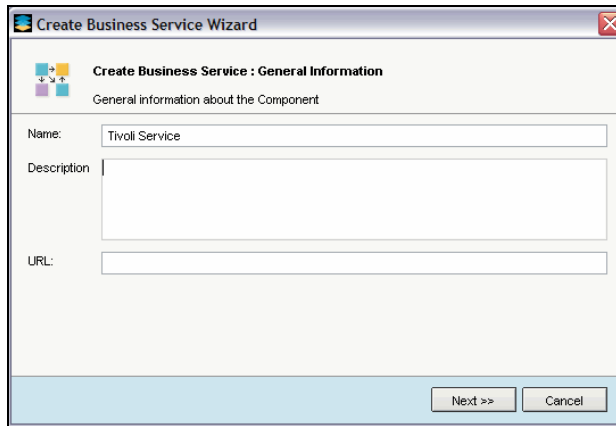
Figure 4-49   Creating a Business Service
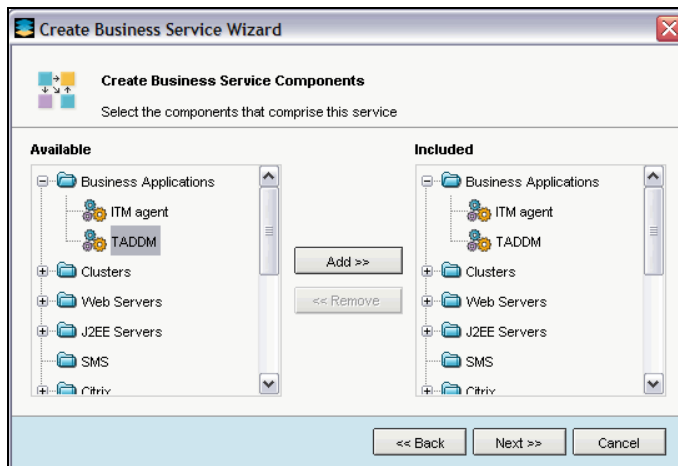
7. Enter a Business Service Name and Description as shown in Figure 4-50. Click **Next**.

Figure 4-50   Creating Business Service Wizard

8. In the next window, you can add Business Applications into the Business Service. We selected **ITM agent** and **TADDM** as shown in Figure 4-51. Click **Next**.

Figure 4-51   Assigning Business Applications to the Business Service

9. Notice that the name of Tivoli Service is added in the Topology tab as a Business Service. Click **Tivoli Service**, and you will see the Business Applications that comprise this service in the right pane as shown in Figure 4-52.
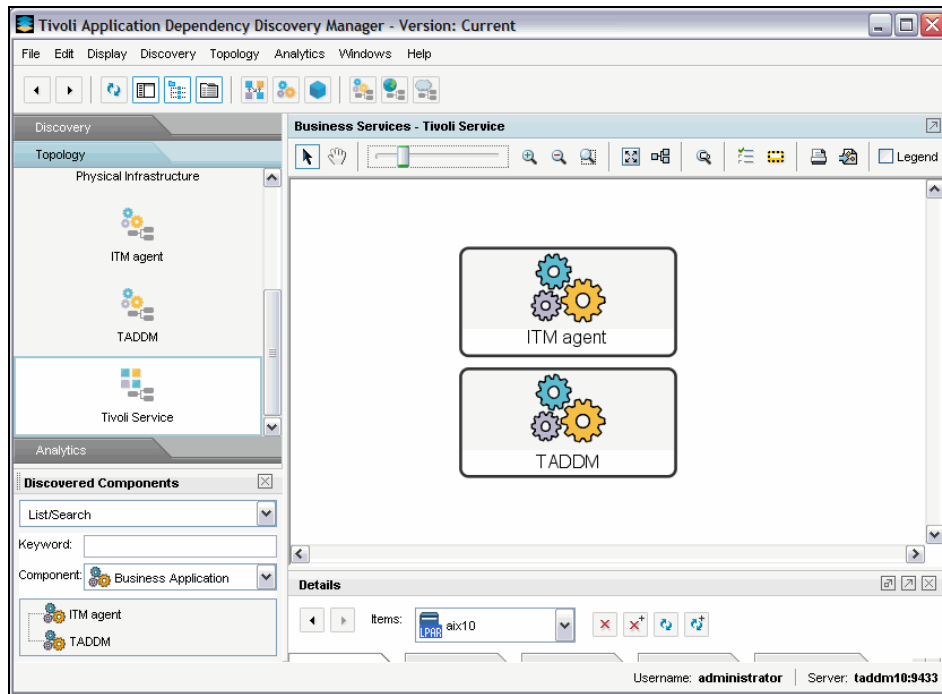
Figure 4-52   Showing the Business Service

10. If you want to see the physical topology of ITM agent, right-click **ITM agent** and select **Show Physical Topology** from the drop-down menu as shown in Figure 4-53.
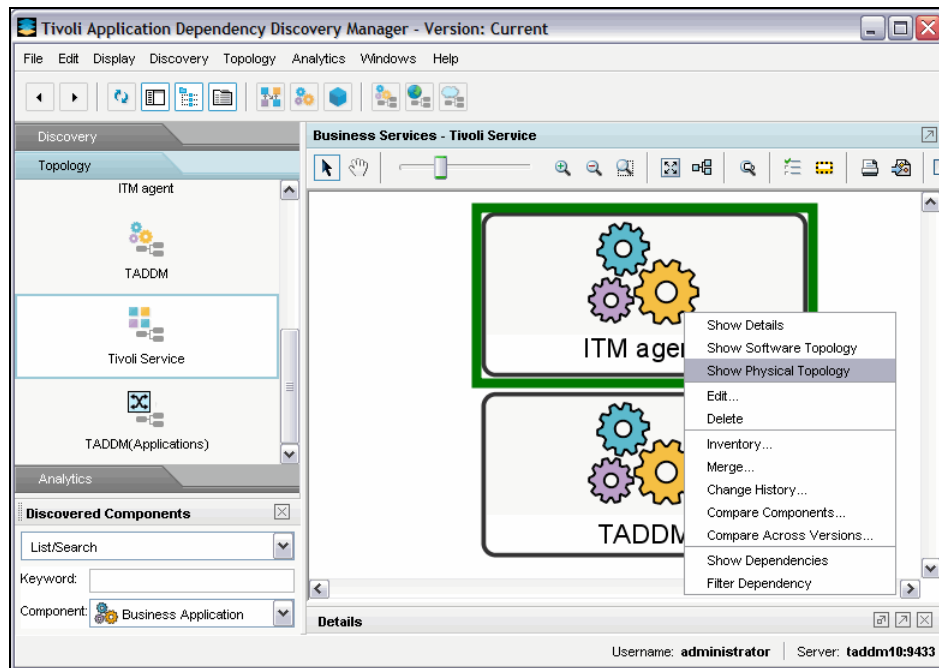
Figure 4-53 Selecting Physical Topology

11.In the next window, the physical topology appears as shown in Figure 4-54.
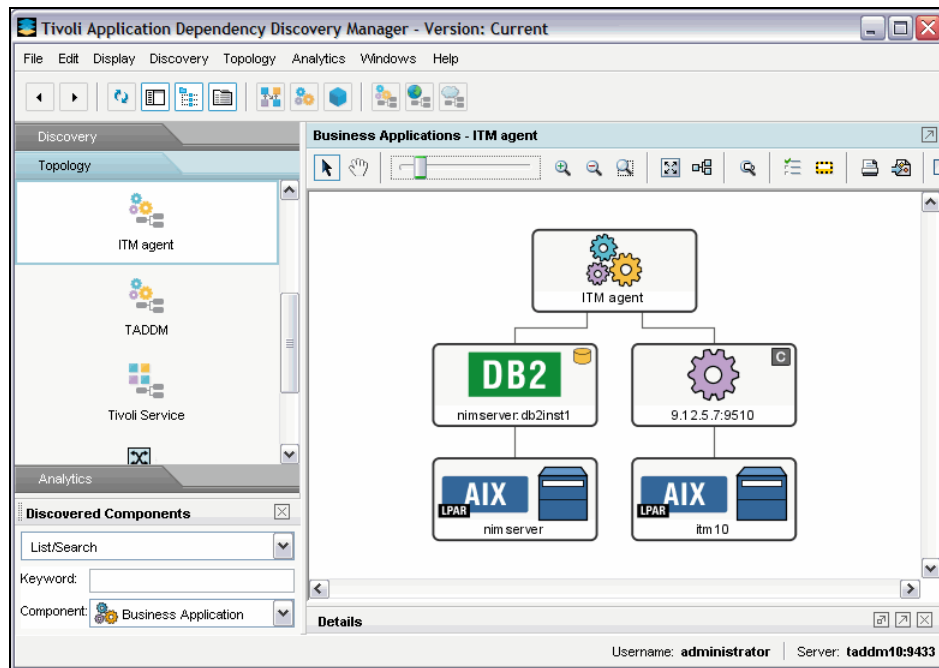
Figure 4-54   Showing the physical topology

4.3.4 Analytics and reporting

You can generate reports and perform analytics to monitor and troubleshoot configuration changes in your environment. You can produce a consolidated summary of changes for a specific Business Application or throughout your entire environment. You can also use the product console to generate an inventory of components that are in your organization.

Change history

If you want to check whether any changes have been made on your applications or servers in a certain period of time, there is a Change History menu under the Analytics tab menu as shown in Figure 4-55 on page 229.

We intentionally add a host, `100.100.100.1 test1` (see the last line), into the `/etc/hosts` file, so that the change can be discovered in TADDM.

Example 4-5   Editing /etc/hosts

```
taddm10:/ ] vi /etc/hosts

9.12.5.38       vio1_l10

9.12.5.32       nimserver   nimserver.itso.ibm.com

9.12.5.118      wpar01      wpar01.itso.ibm.com
9.12.5.119      wpar02      wpar02.itso.ibm.com

9.12.5.36       hmc1        hmc1.itso.ibm.com
9.12.5.33       hmc2        hmc2.itso.ibm.com

100.100.100.1   test1
```

1. Go to the **Discovery** tab menu and click **Run Discover**. When the discovery finishes, come back to the **Change History** menu and click **Run Report**. We set Date Range to `1 hour` as shown in Figure 4-55.
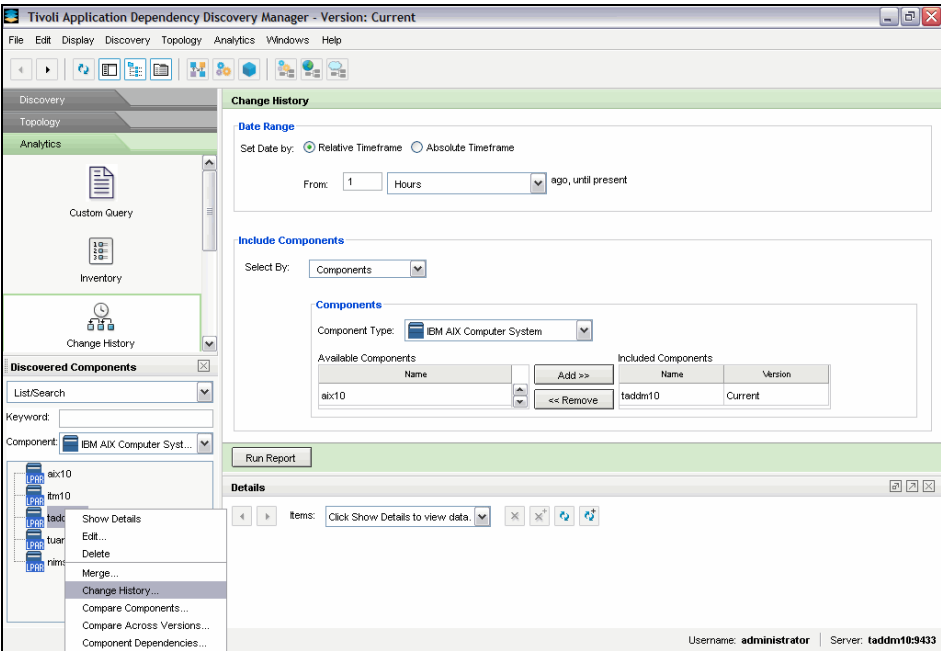
Figure 4-55   Change History

The report shows all the changes that have been made on the server for the past one hour, and you can see the `/etc/hosts` file in the list as shown in Figure 4-56.
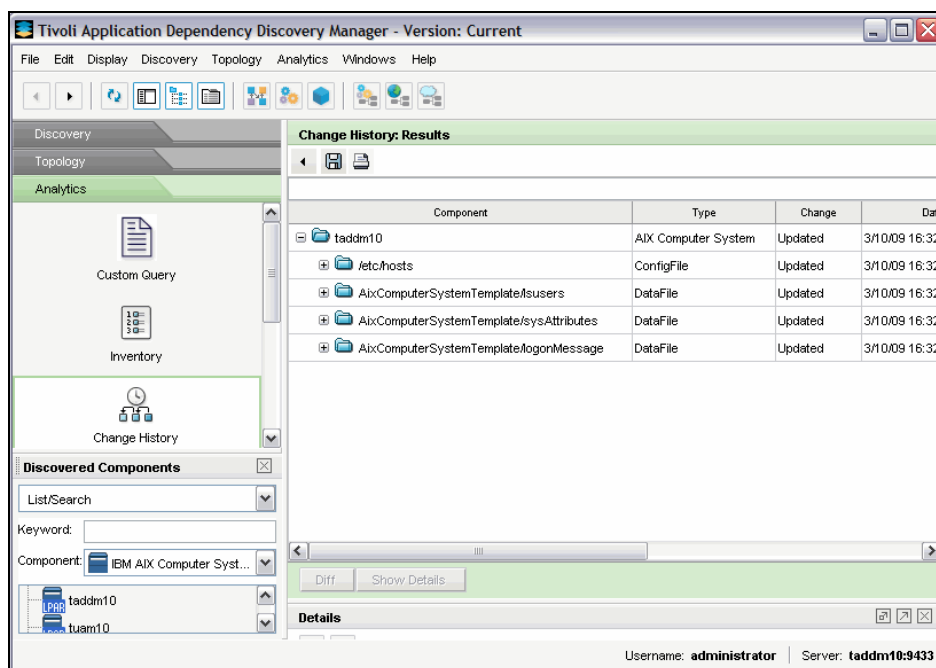
Figure 4-56   Results of Change History

However, note that TADDM does not collect all the files in a server for Change History by default. You might need to check the list in a component and add additional files to be collected.

The next two windows show how to add those files into a computer system.

2. Go to the **Discovery** tab, and then click **Computer Systems**.

3. In the right pane, click **AixComputerSystemTemplate** and then click **Edit** as shown in Figure 4-57.
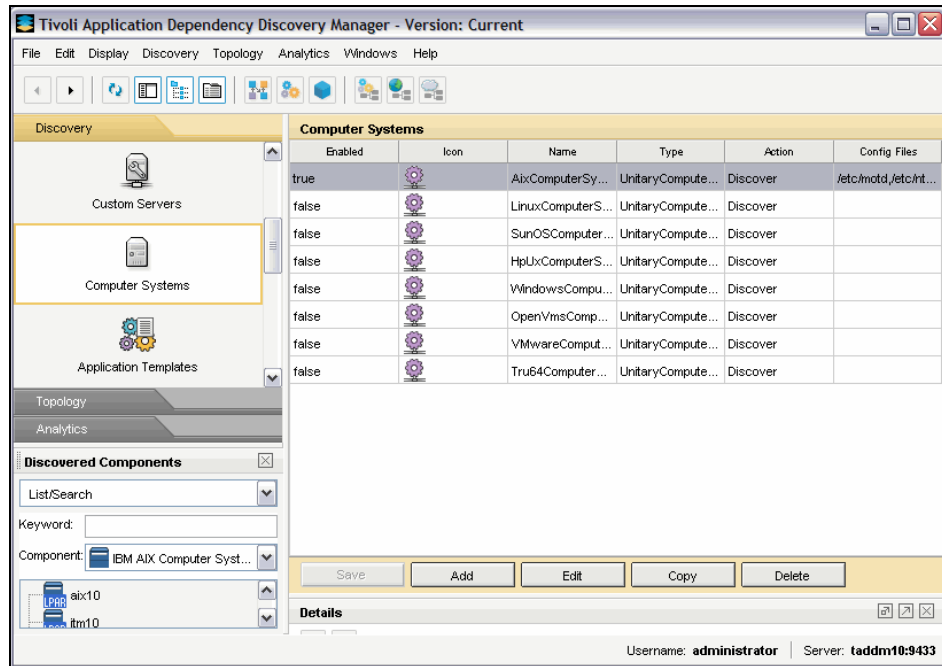
Figure 4-57   Editing Computer Systems

4. In the next window as shown in Figure 4-58, click the **Config Files** tab and add your desired files into the list.
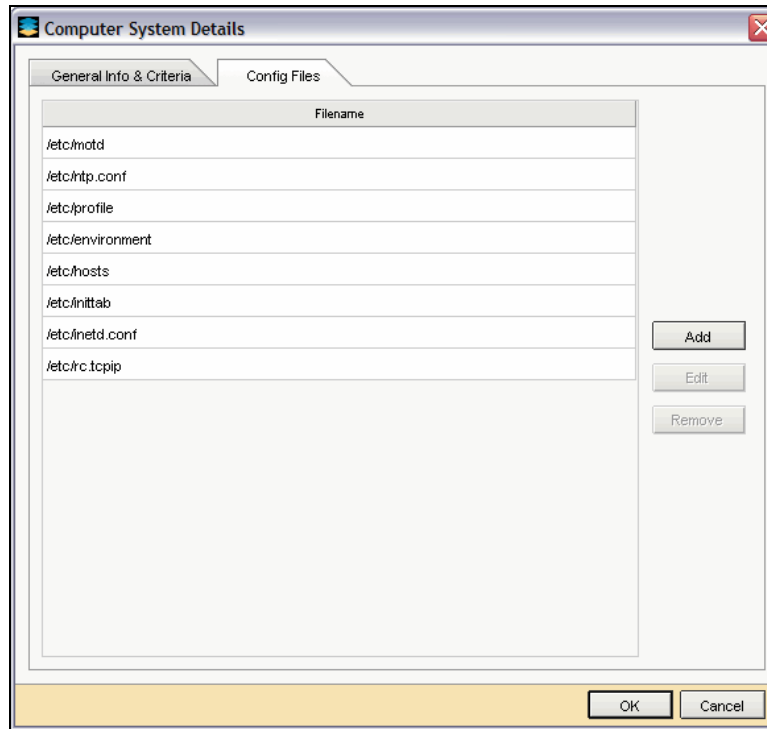
Figure 4-58   Adding Config Files

Inventory reporting

The Inventory report lists components of a specific type with important attributes, for example, it lists all of the AIX Computer systems and their installed OSs and all of the database servers and their versions. It is a comprehensive infrastructure asset report. After you run the inventory report, you can save the data to various formats. If you select multiple types of IT components, after you select the target and file name, multiple files are created.

In order to create a inventory report:

1. Click the **Analytics** tab → **Inventory**, and then select the component type for which you want to create a report as shown in Figure 4-59. We selected **DB2** and **AIX** in this example. Click **Run Report** to create reports.
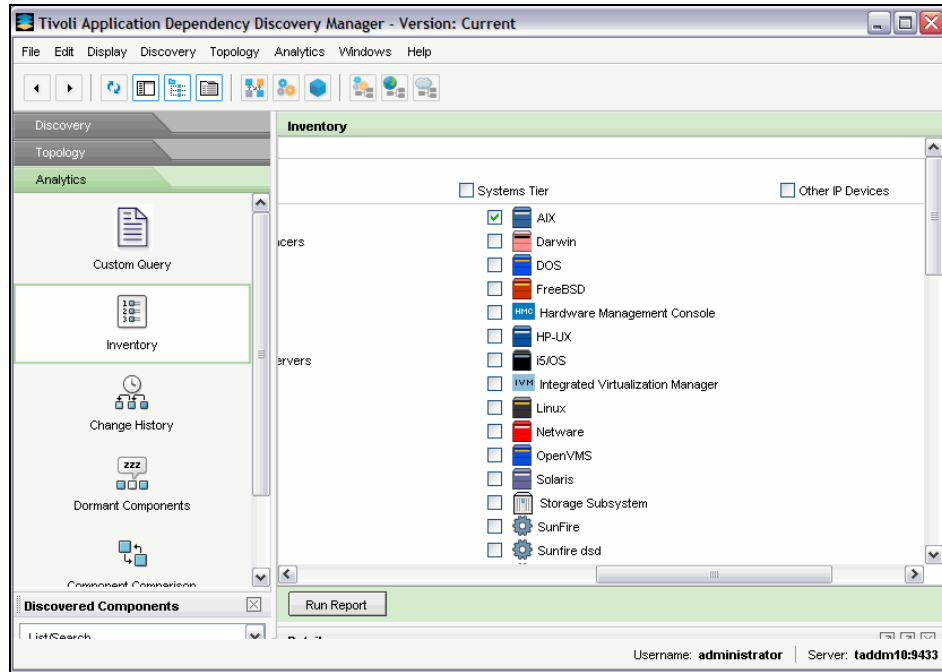
Figure 4-59   Selecting component type for Inventory report

2. Figure 4-60 is an example of a report that includes AIX Servers and the DB2 database. Click the **Save** icon (see cursor) to save the report to your local workstation.
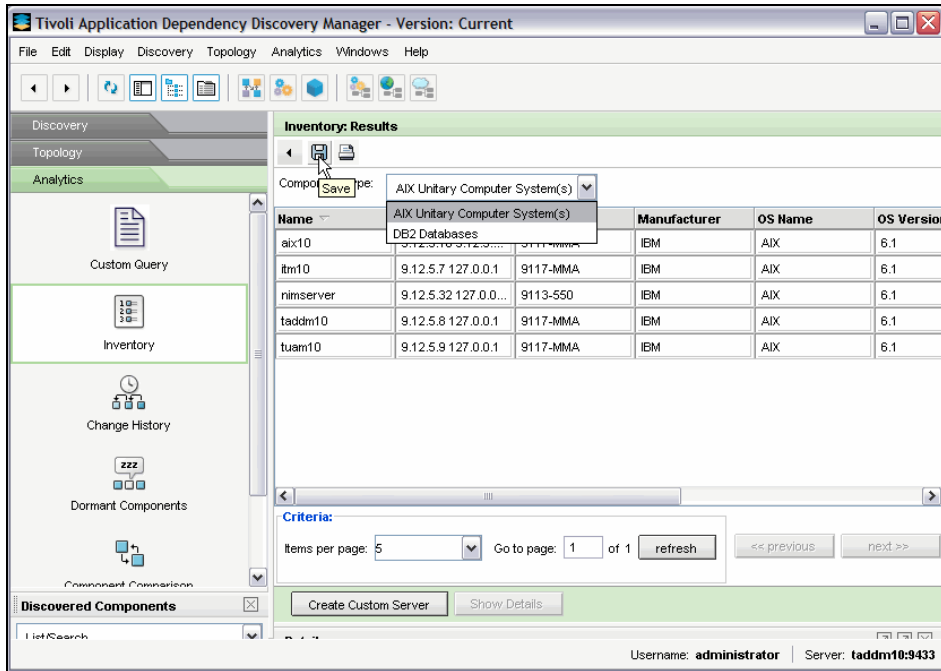
Figure 4-60   Inventory report by component type

3. Provide a file name and select the file type. You can choose various types of files, including portable document format (PDF) or comma-separated values (CSV), as shown in Figure 4-61.
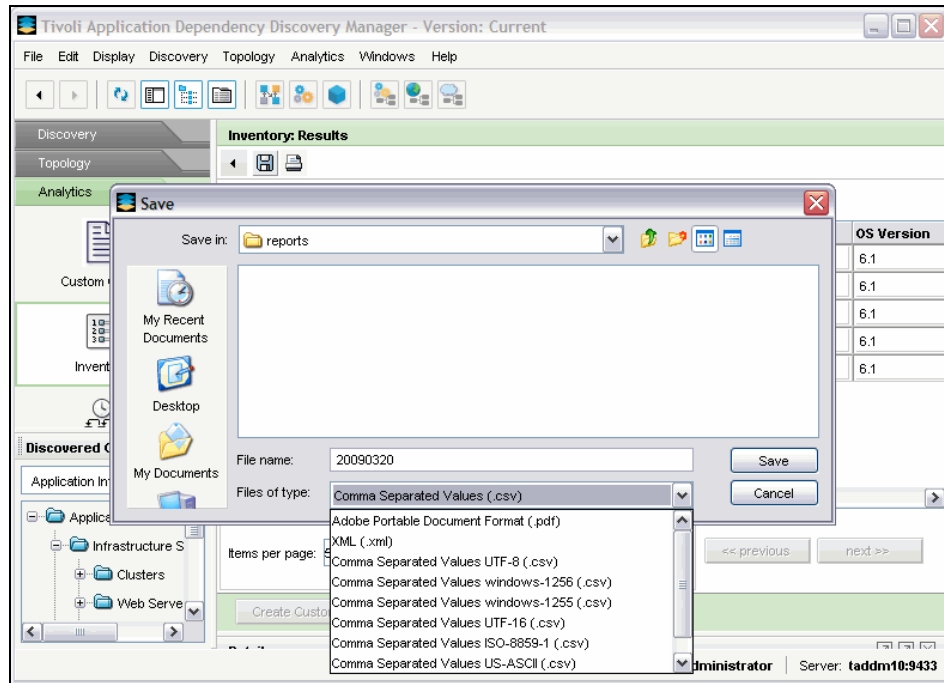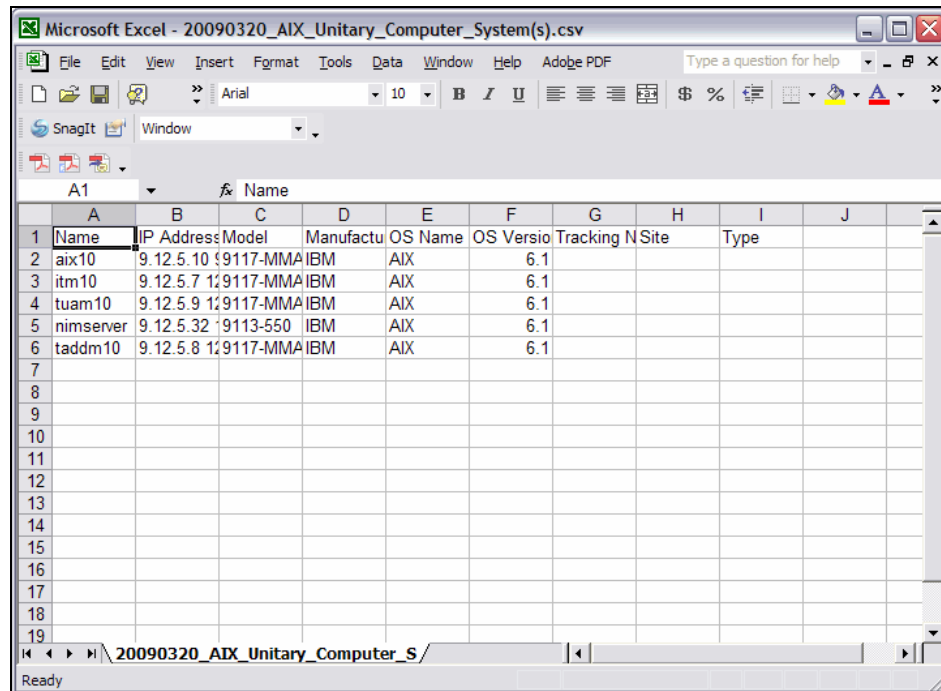
Figure 4-61   File name and type for the inventory report

4. Figure 4-62 shows the sample output of our inventory reports.

Figure 4-62   Sample output of inventory report

If you want to produce reports with Business Intelligence and Reporting Tools (BIRT), refer to Chapter 7, "Reporting Scenarios," in *Deployment Guide Series: IBM Tivoli Application Dependency Discovery Manager V7.1*, SG24-7616, for more details.

4.3.5 Performance considerations

Deploying TADDM in a production environment is an iterative process that involves initial discovery, the refinement of access lists, the refinement of scopes, the refinement of custom servers, the deployment of application descriptors, and rediscovery.

During this process, it might be necessary to review and refine the various settings that impact performance. Although the default settings generally provide adequate performance immediately, you can substantially increase TADDM performance by tailoring the configuration of TADDM to the environment that it manages.

There are several aspects regarding performance with TADDM, including:

► Discovery tuning
► Storage tuning
► Java Virtual Machine tuning
► Database tuning
► Operating system tuning

If you need further information for performance tuning, refer to *Workload Partition Management in IBM AIX Version 6.1*, SG24-7656.

4.4 Operating IBM Tivoli Usage and Accounting Manager

This section describes the basics of operating Tivoli Usage and Accounting Manager:

► "IBM Tivoli Usage and Accounting Manager collectors" on page 237

► "Collecting AIX and VIOS accounting data" on page 241

► "Loading AIX accounting data into the database" on page 242

► "Installing Business Intelligence and Reporting Tools (BIRT)" on page 250

► "Troubleshooting IBM Tivoli Usage and Accounting Manager" on page 266

4.4.1 IBM Tivoli Usage and Accounting Manager collectors

IBM Tivoli Usage and Accounting Manager ships with three major processing scripts as shown in Table 4-3.

Table 4-3 IBM Tivoli Usage and Accounting Manager processing scripts

Script name	Function
/opt/ibm/tivoli/ituam/collectors/Unix/scripts/aacct/ituam_format_aacct	Generates usage logs from data files in the history directory
/opt/ibm/tivoli/ituam/collectors/Unix/scripts/aacct/ituam_get_aacct	Retrieves the *Active* data file and copies it to the history directory; you can run this script multiple times a day if needed
/opt/ibm/tivoli/ituam/collectors/Unix/scripts/aacct/ituam_send_aacct	Sends data files from collector client system to IBM Tivoli Usage and Accounting Manager server

Each of these scripts uses common variables that are set in the `A_config.par` file, which is located in the `/opt/ibm/tivoli/ituam/collectors/Unix/data` directory. The `ituam_format_aacct` and `ituam_get_aacct` scripts do not require special variables to be set up, but in order to use the `ituam_send_aacct` to push the data to the IBM Tivoli Usage and Accounting Manager server, you must set the following variables in the `A_config.par` file:

► CS_METHOD

Method for transferring the data to the IBM Tivoli Usage and Accounting Manager server. This method can be FTP, secure copy protocol (**scp**), or Secure File Transfer Protocol (SFTP).

► CS_PLATFORM

Name of IBM Tivoli Usage and Accounting Manager server.

► CS_USER

Name of user on IBM Tivoli Usage and Accounting Manager server for transfer. We recommend that you set up a specific user for the transfer and that you do *not* use the root user.

► CS_KEY

Password for the user to perform the transfer. If you use **scp**, you can set up keys so that you are not required to have the password in the file in clear text.

► CS_COLL_PATH

Path on the server for the files to be transferred.

► CS_PROC_PATH

Path for processing the data files.

The two processing scripts for gathering the accounting data and formatting the accounting data can be combined into one process. We have included a sample script that runs both of these jobs and that can be used in crontab as a timed process. Example 4-6 shows this script.

Example 4-6   IBM Tivoli Usage and Accounting Manager processing script: tuam-getdata.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 processing needed
else
#date changed, run format script for the previous day
${ITUAMFORMAT} ${START}
fi
```

There are several cronjobs that you add that have been predefined to run the
IBM Tivoli Usage and Accounting Manager processing jobs on the IBM Tivoli
Usage and Accounting Manager server. Example 4-7 lists these jobs.

Example 4-7   IBM Tivoli Usage and Accounting Manager cron jobs

```
# TUAM UNIX/Linux Data Collector scripts
#
5 1 * * * (/opt/ibm/tuam/collectors/Unix/etc/ituam_uc_nightly 1>
/opt/ibm/tuam/collectors/Unix/log/ituam_uc_nightly.log 2>&1)
3,13,23,33,43,53 * * * * /opt/ibm/tuam/collectors/Unix/etc/check_pacct
5 3 * * *
(/opt/ibm/tuam/collectors/Unix/scripts/enterprise/CS_nightly_consolidat
ion 1> /opt/ibm/tuam/collectors/Unix/log/CS_nightly_consolidation.log
2>&1)
#45 3 * * * (/opt/ibm/tuam/collectors/Unix/scripts/enterprise/CS_send
1> /opt/ibm/tuam/collectors/Unix/log/CS_send.log 2>&1)
```

Table 4-4 on page 240 describes each of these jobs.

Table 4-4   IBM Tivoli Usage and Accounting Manager cron job descriptions

Script name	Default execution	Description
ituam_uc_nightly	01:05 a.m.	Collects raw AIX accounting data and file system usage.
check_pacct	Every 10 minutes	Manages `pacct` file.
CS_nightly_consolidation	03:05 a.m.	Consolidates the nightly accounting and storage files into the common source format (CSR) file.
CS_send (optional)	03:45 a.m.	Transfers the consolidated CSR file to the IBM Tivoli Usage and Accounting Manager server.

The scheduled job produce log files that contain information about the execution of each task. The scripts also produce data files that contain accounting and file system information.

Table 4-5 shows the name of the output log and the data file for each of the scheduled cron jobs. The output log directory is `/opt/ibm/tuam/collectors/unix/log`.

Table 4-5   IBM Tivoli Usage and Accounting Manager cron job log files

Script name	Output log file	Output data directory
ituam_uc_nightly	ituam_uc_nightly.log	../collectors/unix/accounting/*<host name>*
check_pacct	-	-
CS_nightly_consolidation	CS_nightly_consolidation.log	../collectors/unix/CS_input_source
CS_send(optional)	CS_send.log	-

4.4.2 Collecting AIX and VIOS accounting data

After the collectors have been installed on all the systems, you need to set up a method to gather the data back to the Tivoli Usage and Accounting Manager server. The two methods of transferring the accounting data to the IBM Tivoli Usage and Accounting Manager server are:

► Push mode, where the client 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 clients

For security reasons, we chose the pull mode for our example. We also chose to implement the pull method using the **scp** command.

We have written a sample script for collecting the AIX accounting data from the nodes, which is shown in Example 4-8.

Example 4-8   Data collection: tuam-getdata.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 36 38 39"

# 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}.t
        xt ${SERVERPATH}/AACCT_${TYPE}/${CLIENT}/
    done
done

# Load data from all the clients
cat $TUAMHOSTS|while read CLIENT

do

echo Loading data from ${CLIENT}

cat ${TUAMHOSTS}/AIXAA.xml|sed "s/sample/${CLIENT}/g" |  sed
"s/\"AIXAA\"/\"AIXAA -${CLIENT}\"/g" > ${TUAMJOBS}/AIXAA-${CLIENT}.xml

chmod 755 ${TUAMJOBS}/AIXAA-${CLIENT}.xml
${JOBRUNNER} ${TUAMJOBS}/AIXAA-${CLIENT}.xml
done
```

> **Note:** You will need to modify the `AIXAA-${CLIENT}.xml` job files that are
> created by the script in Example 4-8 to contain the directories for processing
> your data that is being transferred. The files can be edited by using the
> Integrated Solutions Console (ISC) under **Chargeback Maintenance → Job
> Runner → Job Files**.

4.4.3 Loading AIX accounting data into the database

Tivoli Usage and Accounting Manager jobs process the AIX accounting data
after it has been collected and load it into the Tivoli Usage and Accounting
Manager database.

You can copy the sample job for processing the AIX data and then edit it for your environment:

1. First, create a new job in the ISC. To create a new job, click **Usage and Accounting Manager** → **Chargeback Maintenance** → **Job Runner** → **Job Files** → **New**. Enter the name of your new job file and click **OK.** We want to create a job named `AIXAA.xml`. Refer to Figure 4-63. After the job is created, you will see a message stating, **"`Successfully copied file.`"** This message indicates that the file is created and contains a skeleton job file template.

   Tivoli Usage and Accounting Manager provides a sample job to load AIX data named `SampleAIXAA.xml`. We will use the contents of the `SampleAIXAA.xml` to overwrite the contents of the new `AIXAA.xml` file that we have created.
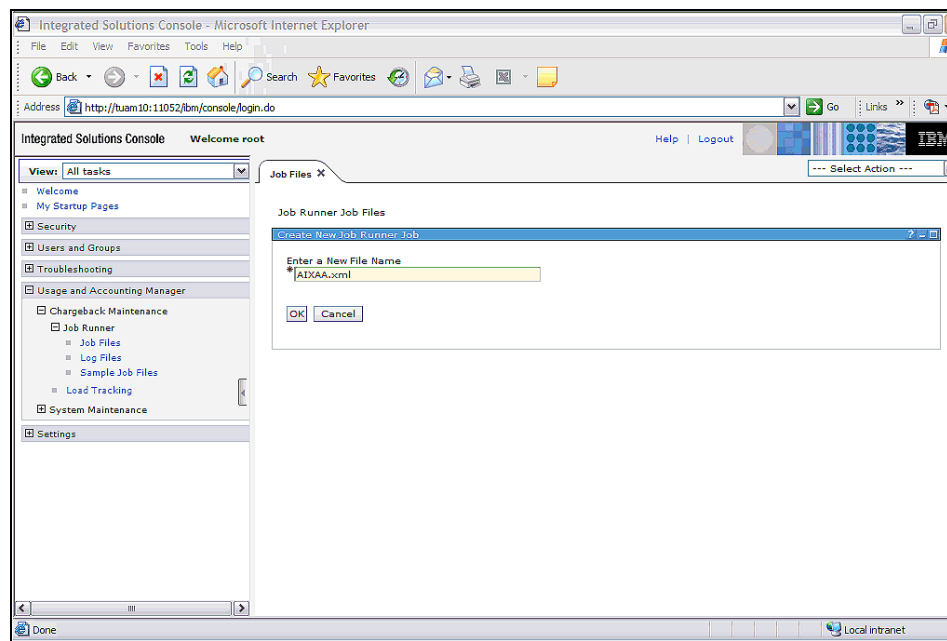
Figure 4-63   Creating a new job by using the ISC

2. From the ISC, click **Sample Job Files** → **SampleAIXAA.xml**. Highlight the contents of SampleAIXAA.xml, then right-click to select the contents of the xml file and click **Copy**. Figure 4-64 on page 244 shows the SampleAIXAA.xml file.

3. From the ISC, click **Usage and Accounting Manager** → **Chargeback Maintenance** → **Job Runner** → **Job Files**. Select **AIXAA.xml** from the menu. Replace AIXAA.xml with the contents that you copied from SampleAIXAA.xml.
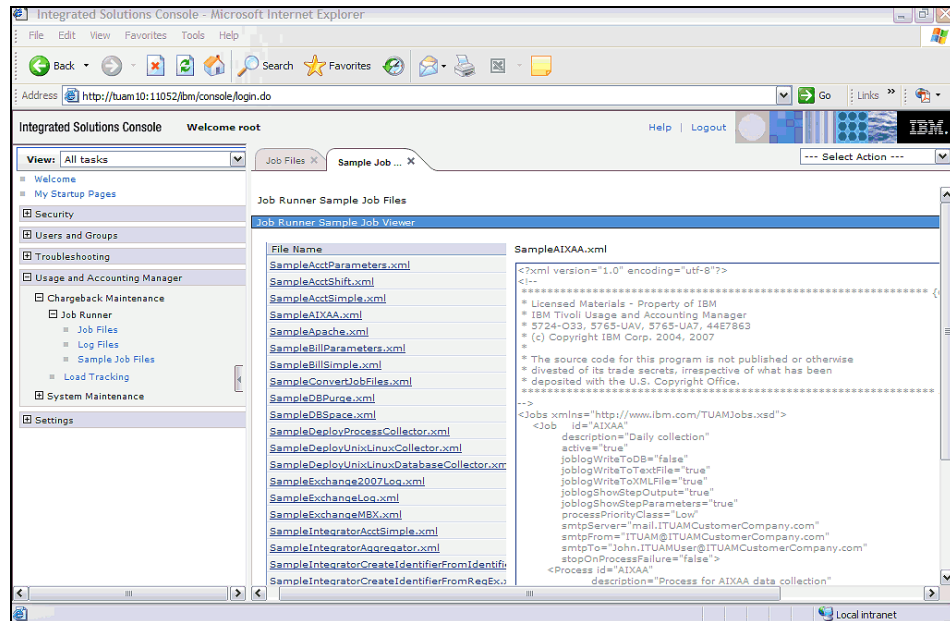
Figure 4-64   SampleAIXAA.xml

The AIXAA.xml file is listed in Example 4-9. After moving over the content of the SampleAIXAA.xml file to the AIXAA.xml file, you need to edit the file to customize it for your environment.

Example 4-9   AIXAA.xml

```
<?xml version="1.0" encoding="utf-8"?>
<!--
  **************************************************************
{COPYRIGHT-TOP}
 * Licensed Materials - Property of IBM
 * IBM Tivoli Usage and Accounting Manager
 * 5724-033, 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="/usr/ituam/CS_input_source/aacct1_%LogDate_End%.txt" />
              <File
name="/usr/ituam/CS_input_source/aacct4_%LogDate_End%.txt" />
              <File
name="/usr/ituam/CS_input_source/aacct6_%LogDate_End%.txt" />
              <File
name="/usr/ituam/CS_input_source/aacct7_%LogDate_End%.txt" />
              <File
name="/usr/ituam/CS_input_source/aacct8_%LogDate_End%.txt" />
```

```
                <File
name="/opt/ibm/tuam/samples/processes/exception.txt" type="exception"
/>
            </Files>
        </Input>

        <Stage name="CreateIdentifierFromIdentifiers" active="true"
trace="false" stopOnStageFailure="true" >
                    <Identifiers>
                        <Identifier name="Account_Code">
                            <FromIdentifiers>
                                <FromIdentifier name="SYSTEM_ID"
offset="1" length="12"/>
                <FromIdentifier name="UserName" offset="1"
length="12"/>
                            </FromIdentifiers>
                         </Identifier>
                    </Identifiers>
                        <Parameters>
                            <Parameter keepLength="true"/>
                            <Parameter modifyIfExists="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="/opt/ibm/tuam/samples/processes/AIXAA/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>
```

4. Run the **AIXAA.xml** job by using the Job Runner in either the ISC or from the command line. In the ISC, click **Usage and Accounting Manager** → **Chargeback Maintenance** → **Job Runner** → **Job Files**. Click **AIXAA.xml** from the list of jobs and select **Run Job** from the pull-down menu, which is shown in Figure 4-65.
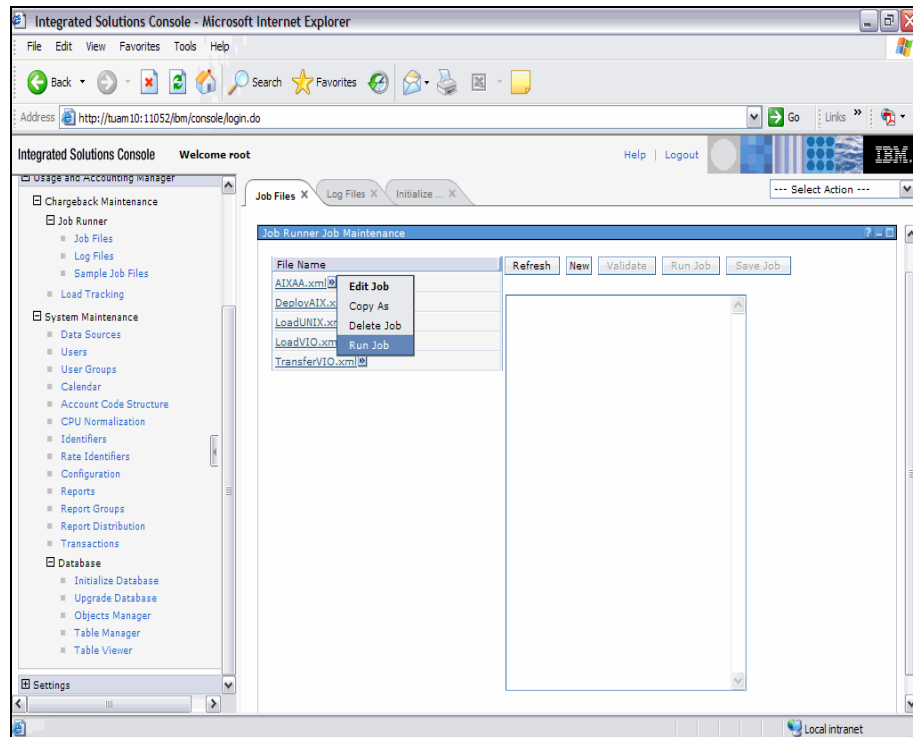
Figure 4-65   IBM Tivoli Usage and Accounting Manager Job Runner

5. When the Job Runner completes, you will be presented with an informational window similar to Figure 4-66. Job Runner can take a few minutes to complete, so wait until this message or a completion message appears.
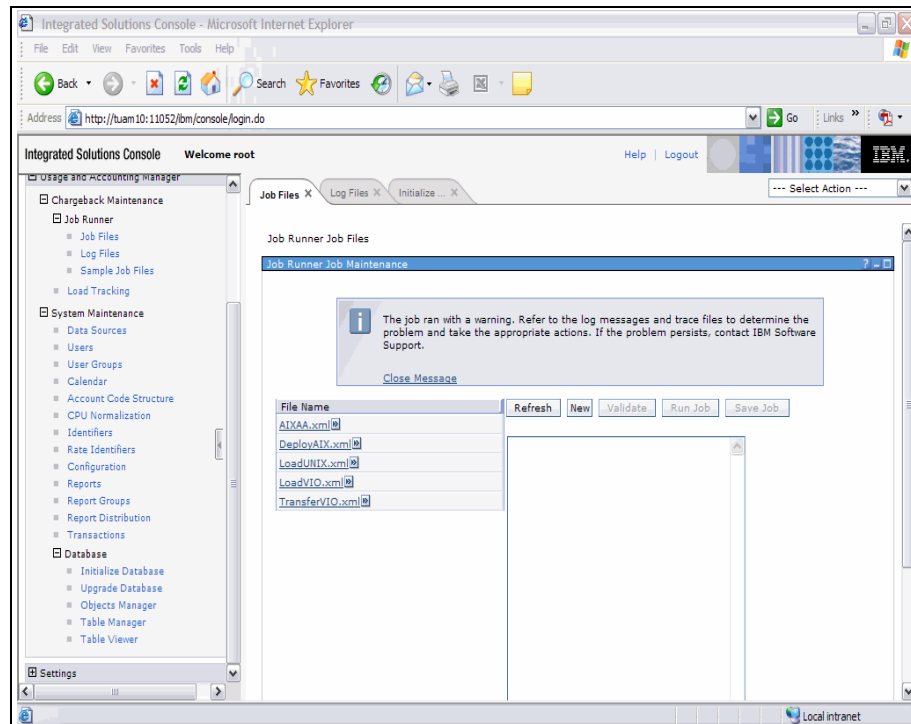
Figure 4-66   IBM Tivoli Usage and Accounting Manager Job Runner complete message

6. Verify that the database has loaded successfully by checking the log files. In the ISC, click **Usage and Accounting Manager** → **Chargeback Maintenance** → **Log Files.** Figure 4-67 shows an example of the log files from the `AIXAA.xml` file. Any errors that occur during the processing will appear in these log files and will be separated into the processing steps. The error messages all begin with `AUC`. You can search for and read the messages on the following Web site:

`http://www-01.ibm.com/software/sysmgmt/products/support/IBMTivoliUsageandAccountingManager.html`
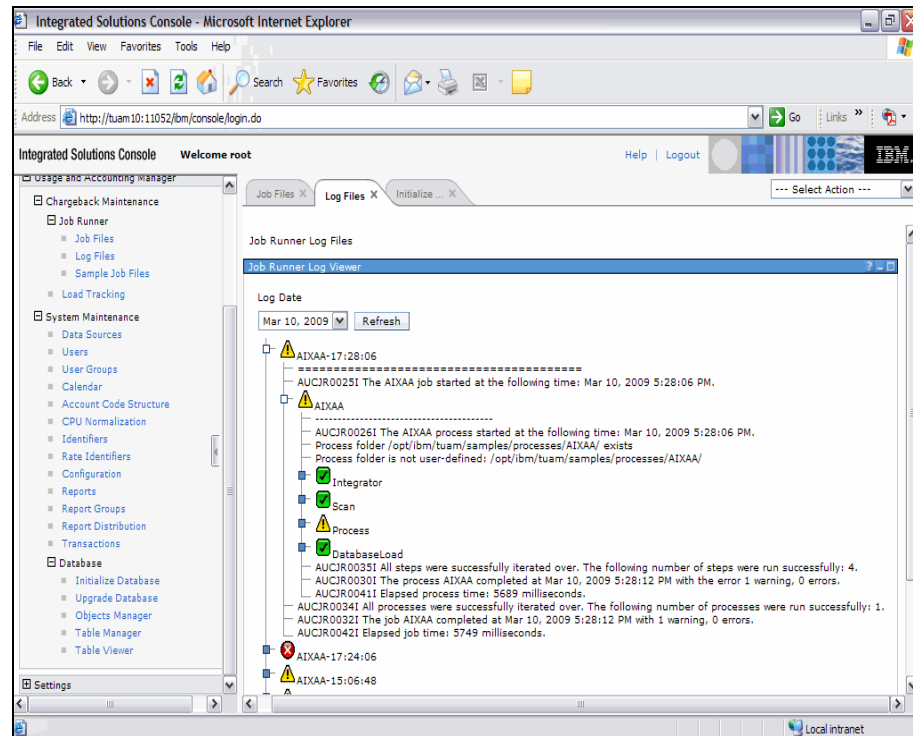
Figure 4-67   Log files

4.4.4 Installing Business Intelligence and Reporting Tools (BIRT)

BIRT is an open source reporting tool that runs in the Eclipse environment. You can choose to download and install BIRT on the Windows, AIX, or Linux platform. BIRT must be run from a graphical user interface (GUI). In our example, we install BIRT on a Windows workstation. Prior to installing and configuring BIRT on a Windows workstation, you must copy the files from the Tivoli Usage and Accounting Manager server that you need to use with BIRT.

These files are located in multiple directories on your IBM Tivoli Usage and Accounting Manager server. The required files are:

- ► `/opt/IBM/db2/V9.1/java/db2jcc_license_cu.jar`

- ► `/opt/IBM/db2/V9.1/java/db2jcc.jar`

- ► `/opt/ibm/tuam/server/reportsbirt/db2/standard/*.rptdesign`

- ► `/opt/ibm/tuam/server/reportsbirt/resources/*`

You must copy these files and place them into a temporary folder for use later. In our example, these files were placed in the `temp` folder:

1. Go to the BIRT Web site at:

   `http://www.eclipse.org/birt/phoenix/`

2. Click **Download BIRT 2.3.2** to go to the BIRT download page.

3. On the BIRT 2.3.2 download page, click **All-in-One for the BIRT Report Engine**. Install Eclipse by unzipping the file.

   Note: If you already have Eclipse installed on your system, you can download the Runtime and Remote Copy (RCP) Designer portions only. These parts are also installed by unzipping the files.

4. Move to the folder that contains the files that you unzipped and look for a folder named *eclipse*. The `eclipse.exe` executable is in the eclipse folder as shown in Figure 4-68. Double-click **eclipse.exe** to start the eclipse environment.
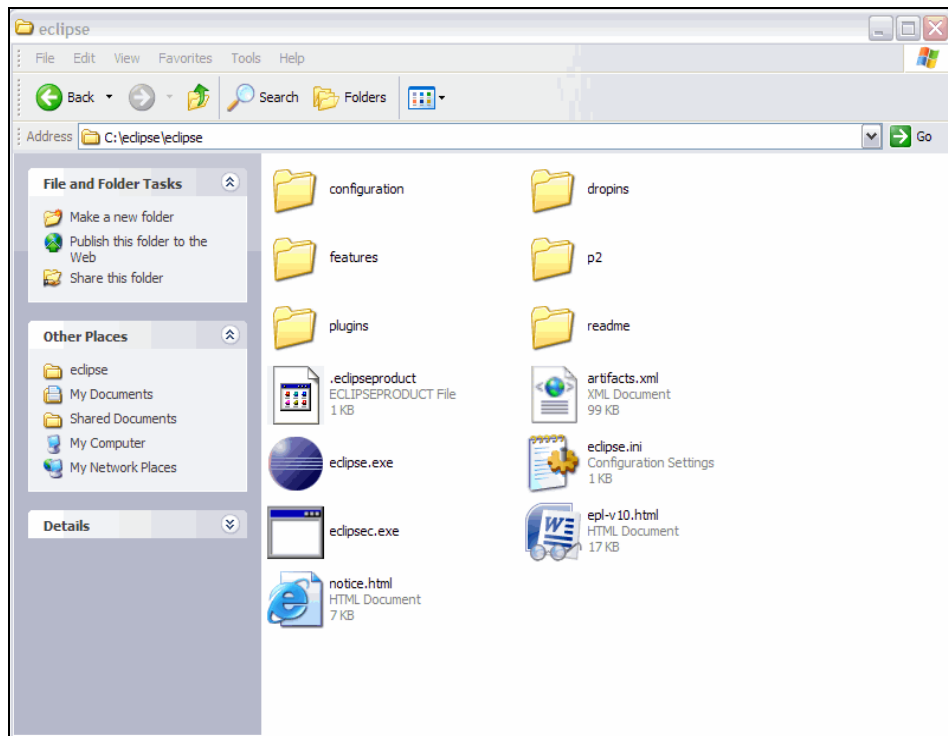
Figure 4-68   Eclipse folder

5. The first time that you execute `eclipse.exe`, you will see a warning message as shown in Figure 4-69. Click **Run** to continue, and also clear the check mark from "Always ask before opening this file."
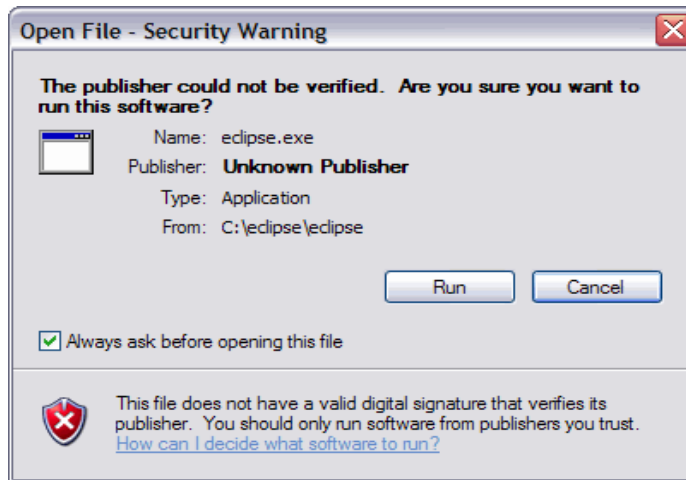
Figure 4-69   Eclipse security warning

6. The Eclipse Workspace Launcher will start and prompt you to enter the path to use for projects as shown in Figure 4-70. Select the path that you will use to save your BIRT projects and select **OK** to continue.

Figure 4-70   BIRT Workspace Launcher

7. BIRT will open up to the main workspace window as shown in Figure Figure 4-71. You need to create a new project by selecting **File → New → Project**. A pop-up window appears as in Figure 4-71. Expand **Business Intelligence and Reporting Tools**, select **Report Project**, and click **Next**.
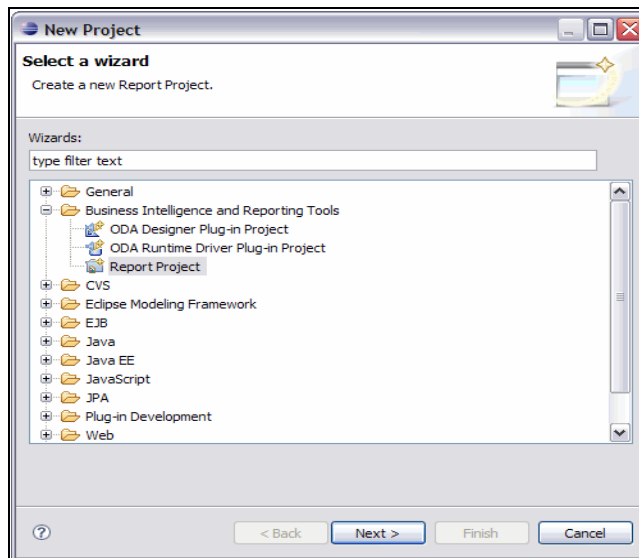
Figure 4-71   New project window

8. You will be prompted for a name for the new project as in Figure 4-72. Enter the new project name, and click **Finish** to continue.
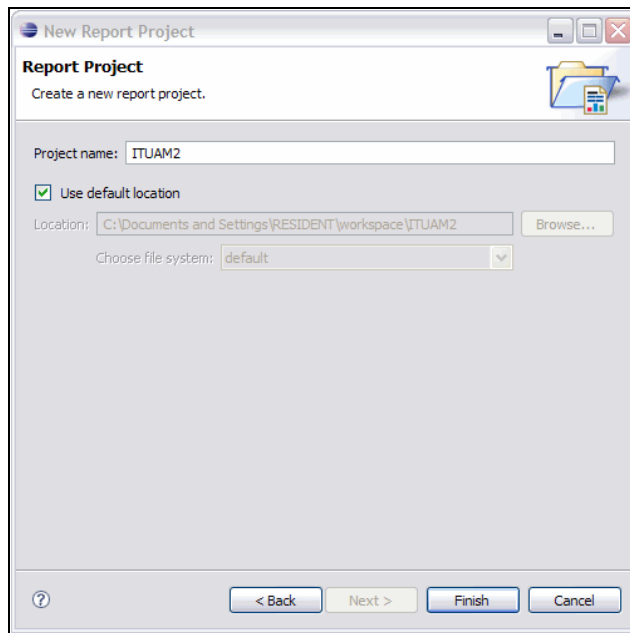
Figure 4-72   New Report Project

9. You will be presented with a blank Report Design workspace similar to the one in Figure 4-73.
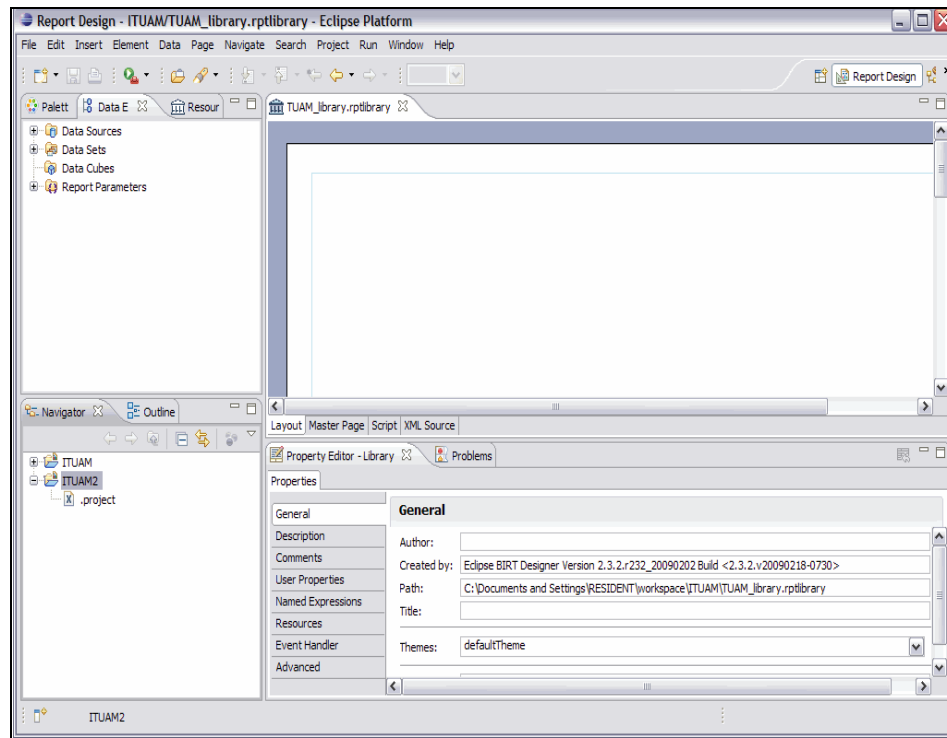
Figure 4-73   BIRT main window

Configuring BIRT

In order for BIRT to connect to the DB2 server on the Tivoli Usage and Accounting Manager server, we need to configure Java Database Connectivity (JDBC) drivers and ensure that the data sets are defined with the correct user name and password. To configure BIRT:

1. From the BIRT main window, click the **Data Explorer** tab → **Data Sources**. Right-click **DB2**, and select **Edit** from the pull-down menu. The menu in Figure 4-74 on page 257 appears.
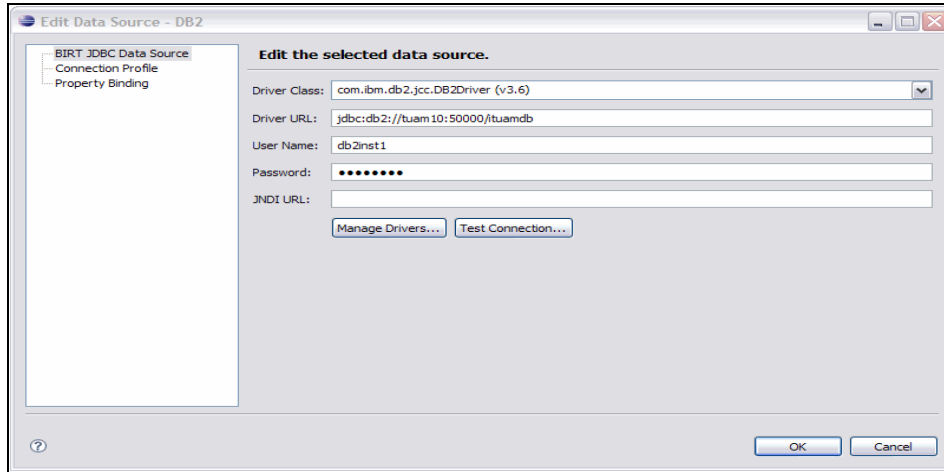
Figure 4-74   Edit data Source

2. Change the Driver URL to `jdbc:db2://`*`yourserver`*`:50000/`*`yourdbname`*, change the user name to `db2inst1`, and change the password to the password that you set for `db2inst1` on your system. Click **Manage Drivers**. The Manage JDBC Drivers window appears as in Figure 4-75.
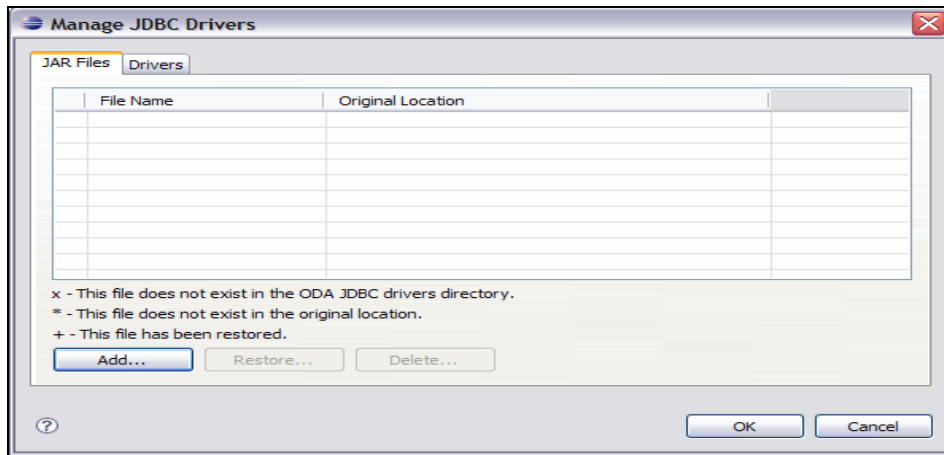
Figure 4-75   Manage JDBC Drivers

3. From this window, you need to add the JDBC drivers to connect to the database.Click **Add** and navigate to the `Temp` folder where you stored the driver files earlier. Select the **db2jcc.jar** file and click **Open**, and repeat for the `db2jcc_license_cu.jar` file. When both files appear as in Figure 4-76 on page 258, click **OK** to continue.
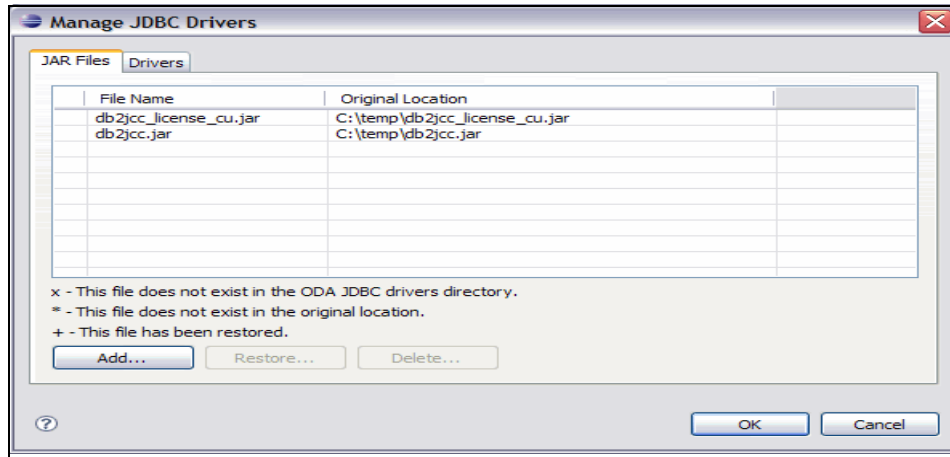
Figure 4-76 Manage JDBC Drivers completion

4. Test the connection to the database as shown in Figure 4-77.

Figure 4-77 Database connection test

Edit SQL statements

The SQL statements that are included in the BIRT reporting methods are written using `dbadmin` as a user ID. We need to edit the statements and put in the correct user ID of `db2inst1` for the statements to work correctly with the DB2 environment. We also need to ensure that the Data Source type for all the data sets is set to DB2. To edit the data sets in BIRT:

1. From the main window in the Data Explorer Tab, click **Data Sets**. Right-click the name of the Data Set (for example, `DB2_GET_CLIENT`), and select **EDIT** from the drop-down menu.

2. The window in Figure 4-78 appears.

3. Change the Data Set type to **DB2**.
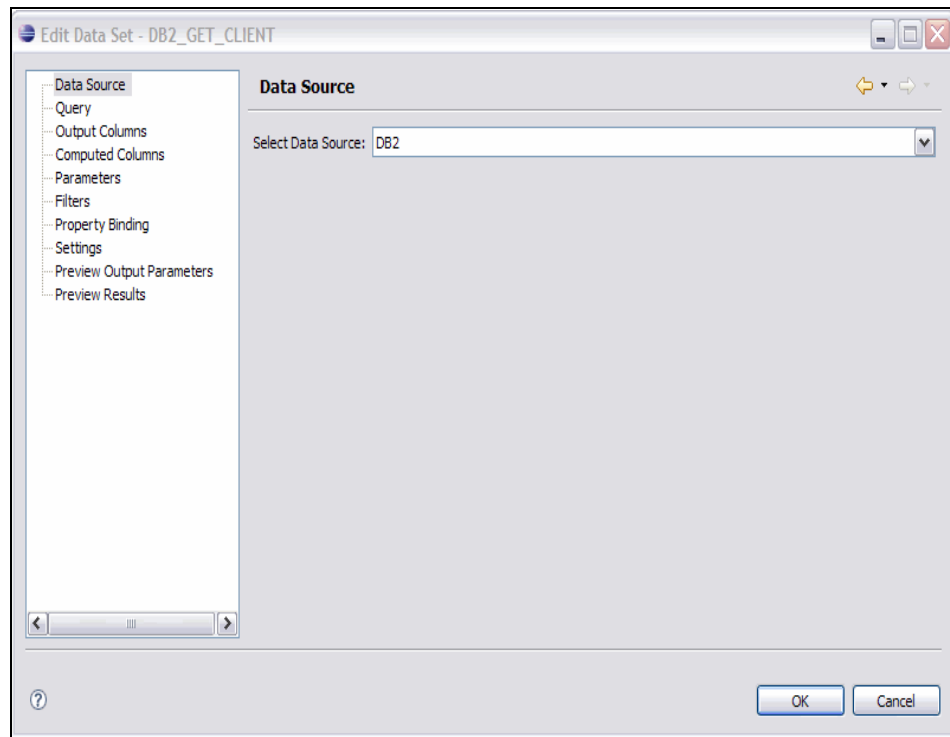
Figure 4-78   Edit Data Set

4. Next, click the Query line to edit the Data Set query. Change the user name that is in the query to `db2inst1`. The line that contains the user name is `{call DBADMIN.GET_CLIENT(?,?)}`. Change `DBADMIN` to `DB2INST1` as shown in Figure 4-79. Click **OK** to continue. Complete the same process for all the data sets, including the SQL data sets, and change the user name in the SQL data sets from `DB0` to `DB2INST1`.
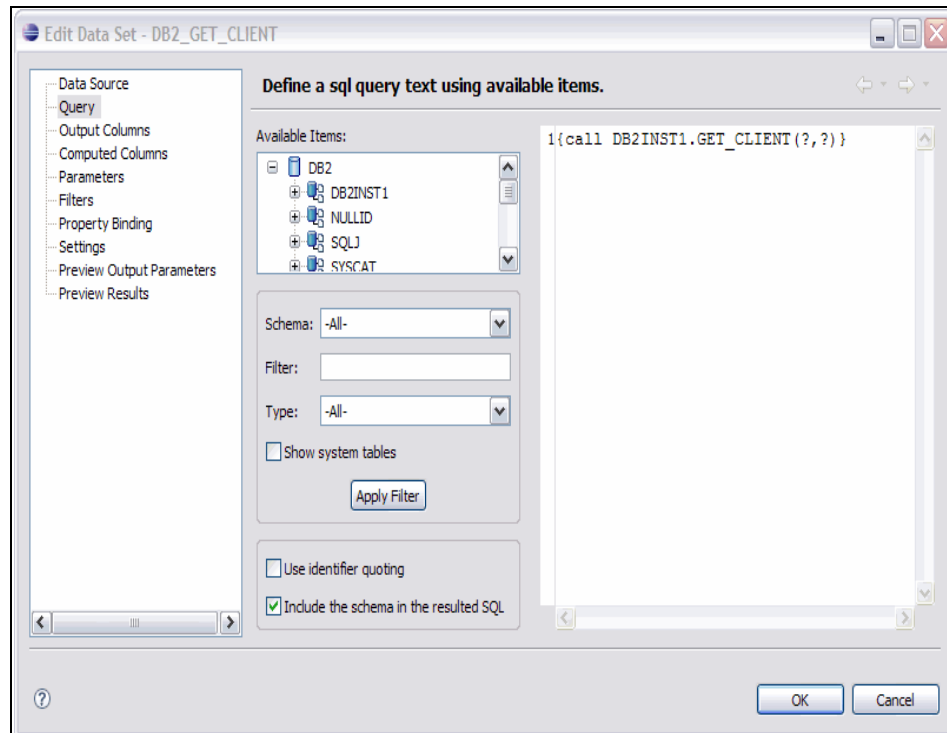
Figure 4-79   Edit SQL Query

Included BIRT reports

Tivoli Usage and Accounting Manager includes five templates for BIRT reporting. The templates are in the `/opt/ibm/tuam/server/reportsbirt/db2/standard` directory. The templates are:

`RCFGX001.rptdesign`   Configuration Report

`RCLTX001.rptdesign`   Client Report

`RINVC007.rptdesign`   Invoice

`RIVTC001.rptdesign`   Run Total Invoice

`RRATX001.rptdesign`   Rate Report

Sample reports

You can use the BIRT templates to generate reports for date ranges from your IBM Tivoli Usage and Accounting Manager database:

1. To run a report in the eclipse project Navigator window, right-click the report type that you want from the list. Select **Report** → **Run Report**. Figure 4-80 shows the parameters that you enter for a Run Total Invoice report.
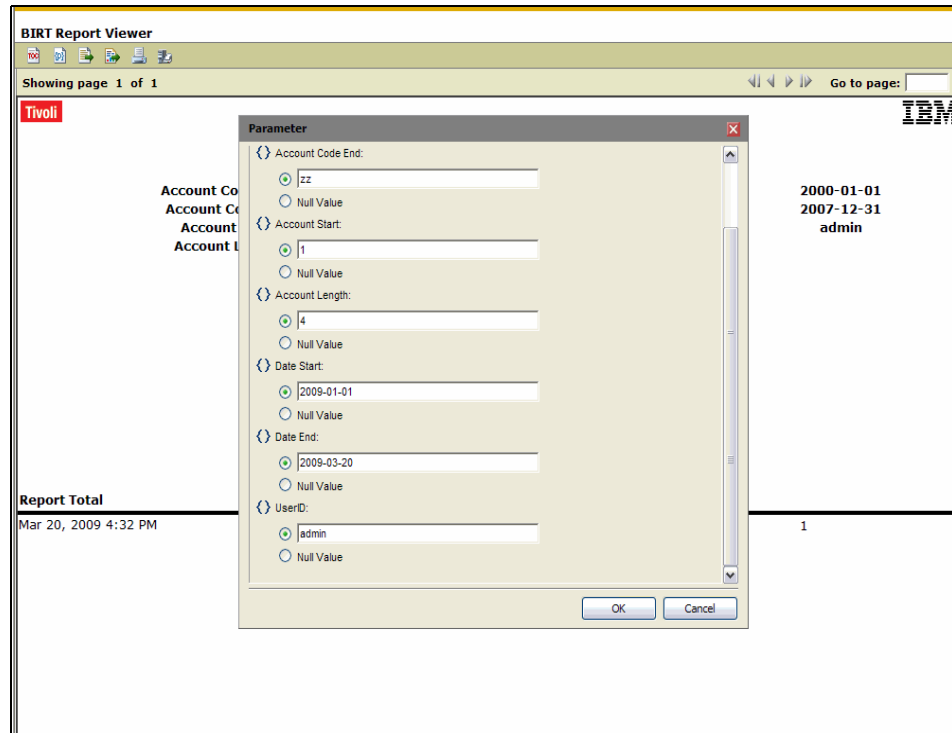
Figure 4-80   Setting report parameters

2. After you set the report parameters, you will see your data on the report for the dates that you have selected. Figure 4-81 shows a sample Run Total Invoice report from our test environment.

Figure 4-81   Run Total Invoice BIRT report

3. BIRT reports can be exported into several formats, including Microsoft® Excel®, postscript, Microsoft Word, Microsoft PowerPoint®, and PDF. Figure 4-82 on page 264 shows the PDF of the Run Total Invoice report. You can use these formats to present reports to customers about system usage and relate the system usage back to actual accounting charges.

Figure 4-82   BIRT Run Total Invoice report PDF

An invoice report can be seen in Figure 4-82 on page 264.

Figure 4-83   BIRT Invoice

Figure 4-84 shows the BIRT Rate Report that is available in the templates.

Figure 4-84   BIRT Rate Report

For additional information about producing reports with BIRT, refer to Chapter 7, "Reporting Scenarios" in the *Deployment Guide Series: IBM Tivoli Application Dependency Discovery Manager V7.1,* SG24-7616.

4.4.5 Troubleshooting IBM Tivoli Usage and Accounting Manager

This section describes basic problems that have been encountered running IBM Tivoli Usage and Accounting Manager.

DB2 transaction logs full

You might get the error message "`AUCPE0202E The DBLoad process completed unsuccessfully with the following exception:`" when the DB2 transaction logs are full. The solution is shown in Example 4-10 on page 267.

Example 4-10   Increasing UDB Transaction Logs

```
tuam10:]su - db2inst1
$db2 connect to ituamdb

   Database Connection Information

 Database server        = DB2/AIX64 9.1.4
 SQL authorization ID   = DB2INST1
 Local database alias   = ITUAMDB


$db2 get db cfg | grep -i LOGSECOND
Number of secondary log files            (LOGSECOND) = 4
$db2 update db cfg using logsecond 150 immediate
DB20000I  The UPDATE DATABASE CONFIGURATION command completed
successfully.
$db2 get db cfg show detail | grep -i LOGSECOND
Number of secondary log files            (LOGSECOND) = 150 150
$db2 get snapshot for db on ITUAMDB | grep -i SEC
Secondary connects total                 = 3
Total buffer pool read time (milliseconds) = Not Collected
Total buffer pool write time (milliseconds)= Not Collected
Maximum secondary log space used (Bytes)  = 208181636
Secondary logs allocated currently        = 50
Log read time (sec.ns)                     = 0.000000004
Log write time (sec.ns)                    = 142.000000004
Application section lookups                = 3429
Application section inserts                = 843
 Total shared section inserts             = 0
 Total shared section lookups             = 0
 Total private section inserts            = 843
 Total private section lookups            = 903
        Secondary ID                         = 2
        Secondary ID                         = 1
        Secondary ID                         = System 32k buffer pool
        Secondary ID                         = System 16k buffer pool
        Secondary ID                         = System 8k buffer pool
        Secondary ID                         = System 4k buffer pool
```

File system filling up

Watch the file system sizes for a full condition as shown in Example 4-11 on
page 268. Increase the file system size or clean out the old accounting
transaction logs. The sample jobs that are set up in this book are set to delete the

accounting files after 45 days. Be aware that the data is stored in multiple locations.

The first location where the data is stored is on the client server in `/opt/ibm/tuam/collectors/unix/CS_input_source`, which holds the daily accounting data that is collected on the client. The data is transferred to the server, but a copy is preserved on the client also. Setting up jobs to remove the accounting files periodically is a good practice to prevent "`File system full`" conditions (Example 4-11) from occurring in your environment.

The second location that contains the accounting data is on the server. The accounting data from all the Tivoli Usage and Accounting Manager clients is gathered and stored in both the `/usr/ituam/CS_input_source` directory and then loaded into the database also. Both the database and the file system that contains the raw data for processing will need to be monitored for size.

> **Note**: The `/usr/ituam/CS_input_source` directory is a link to `/opt/ibm/tuam/collectors/unix/CS_input_source`. It is good practice to create this path as a file system when you are collecting data from a large number of clients.

Example 4-11   File system sizes in AIX and file system full condition

```
$ df -k
Filesystem    1024-blocks     Free %Used    Iused %Iused Mounted on
/dev/hd4            49152     2636   95%     1942    69% /
/dev/hd2          2342912    76116   97%    42773    63% /usr
/dev/hd9var        262144   122884   54%     2419     9% /var
/dev/hd3           507904   479236    6%      310     1% /tmp
/dev/hd1            16384    10740   35%       77     4% /home
/dev/hd11admin     131072   130708    1%        5     1% /admin
/proc                   -        -    -        -      -  /proc
/dev/hd10opt       147456    11324   93%     2544    48% /opt
/dev/livedump      262144   261776    1%        4     1%
/var/adm/ras/livedump
/dev/fslv00       2097152  1854484   12%       44     1% /logs
/dev/fslv01       2097152  1253384   41%      448     1%
/home/db2inst1
/dev/fslv02        524288    33484   94%     3756    25% /opt/IBM/db2
/dev/fslv03       1114112     7316  100%    11994    73% /opt/ibm/tuam
/dev/itm_lv       4194304  3753964   11%     3288     1% /opt/IBM/ITM
```

Installation problem with JDBC driver

In the process of configuring IBM Tivoli Usage and Accounting Manager 7.1, you might see the following message:

```
The following information was provided:
[ibm][db2][jcc][10333][11649] No license was found. An appropriate
license file db2jcc_license_*.jar must be provided in the CLASSPATH
setting. Review the trace log to get detailed information.
```

To resolve this problem, you need to add the full path pointing to the location for the DB2 JDBC driver license file and then restart the IBM Tivoli Usage and Accounting Manager server.

To avoid the problem, follow the steps listed in the next section.

Adding a JDBC driver

The appropriate JDBC driver must be available to enable the Usage and Accounting Manager Processing Engine to access the DB2 database. If a client needs to connect to the database, you are required to install the driver for that client.

To enable the Usage and Accounting Manager Processing Engine to access the database, the appropriate JDBC driver must be available on the server running Tivoli Usage and Accounting Manager. The DB2 `db2jcc.jar` and `db2jcc_license_cu.jar` (license JAR file) drivers are used by Tivoli Usage and Accounting Manager. If you select JDBC support during the installation steps for DB2, the driver and the license file are included in the installation.

To locate the drivers:

1. In the ISC, click **Usage and Accounting Manager** → **System Maintenance** → **Configuration** → **Drivers**.

2. On the Configuration - Drivers page, click **New**. The Browse button and text box become available.

3. Type the path to the driver in the text box or click **Browse** to locate the path to the driver. For DB2 for UNIX, enter the paths to both `db2jcc_license_cu.jar` and `db2jcc.jar`.

4. Click **OK** to add the driver. The driver is added to the list of drivers.

Debugger trace and log information

Tivoli Usage and Accounting Manager has several log files that it uses for the administration server, processing server, reporting server, and database server, as shown in Table 4-6 on page 270.

Table 4-6   Tivoli Usage and Accounting Manager trace and log file information

Path or filename	function
/opt/ibm/tuam/logs/jobrunner/<*JobId*>	Job runner log files separated per JobID
<*time stamp*>.txt	Job log output
<*time stamp*>.xml	Job log for use with ISC
/opt/ibm/tuam/logs/server/	Tivoli Usage and Accounting Manager trace and log files
message0.log message<*g*>.log<*#*>	Messages from tuam processing, where <*g*> = generation and <*#*> = instance
trace0.log trace<*g*>.log<*#*>	Trace details for tuam processing, where <*g*> = generation and <*#*> = instance
*.lck	Lock files for trace and log coordination
/opt/ibm/tuam/ewas/profiles/AppSrv01SystemOut.log/logs/server1	WebSphere and ISC files
SystemOut.log	WebSphere messages
SystemErr.log	WebSphere error log
/opt/ibm/tuam/config	Tivoli Usage and Accounting Manager config files
logging.properties	Settings for trace and message files
jdk_logging.properties	Not used with version 7.1
/opt/ibm/tivoli/common/AUC/logs/install	Installation and uninstallation log files

IBM Tivoli Usage and Accounting Manager additional resources

Additional resources that are useful for continued reading:

► *IBM Tivoli Usage Accounting Manager V7.1 Handbook*, SG24-7404
► *Deployment Guide Series: IBM Tivoli Usage and Accounting Manager V7.1*, SG24-7569

Abbreviations and acronyms

AACCT	AIX Advanced Accounting	**CI**	Configuration items
ABI	Application Binary Interface	**CID**	Configuration ID
ACL	Access Control List	**CLDR**	Common Locale Data Repository
AFPA	Adaptive Fast Path Architecture	**CLI**	Command-Line Interface
AIO	Asynchronous I/O	**CLVM**	Concurrent LVM
AIX	Advanced Interactive Executive	**CCMDB**	Change and Configuration Management Database
APAR	Authorized Program Analysis Report	**CPU**	Central Processing Unit
API	Application Programming Interface	**CSM**	Cluster Systems Management
ARP	Address Resolution Protocol	**CSR**	Common source format
ASMI	Advanced System Management Interface	**CT**	Component Trace
BFF	Backup File Format	**DAS**	Database Administrator Server
BIND	Berkeley Internet Name Domain	**DB**	Database
BIRT	Business Intelligence Reporting Tools	**DCEM**	Distributed Command Execution Manager
BIST	Built-In Self-Test	**DCM**	Dual Chip Module
BLV	Boot Logical Volume	**DES**	Data Encryption Standard
BOOTP	Boot Protocol	**DGD**	Dead Gateway Detection
BOS	Base Operating System	**DHCP**	Dynamic Host Configuration Protocol
BSD	Berkeley Software Distribution	**DLPAR**	Dynamic LPAR
CA	Certificate Authority	**DMA**	Direct Memory Access
CATE	Certified Advanced Technical Expert	**DNS**	Domain Name Server
CD	Compact Disk	**DR**	Dynamic Reconfiguration
CDE	Common Desktop Environment	**DRM**	Dynamic Reconfiguration Manager
CEC	Common Event Console	**DST**	Daylight Saving Time
CHRP	Common Hardware Reference Platform	**eCMDB**	Enterprise Configuration Management Database
		EC	EtherChannel

271

ECC	Error Checking and Correcting	**I/O**	Input/Output
EGID	Effective Group ID	**IBM**	International Business Machines
EJB™	Enterprise Java Bean	**ICU**	International Components for Unicode
EIF	Event Integration Base		
EPOW	Environmental and Power Warning	**ID**	Identification
		IDE	Integrated Device Electronics
EPS	Effective Privilege Set	**IEEE**	Institute of Electrical and Electronics Engineers
ERRM	Event Response Resource Manager		
		IETF	Internet Engineering Task Force
ESE	Enterprise Server Edition		
ESS	Enterprise Storage Server®	**IP**	Internetwork Protocol
EUID	Effective User ID	**IPAT**	IP Address Takeover
F/C	Feature Code	**IPL**	Initial Program Load
FC	Fibre Channel	**IPMP**	IP Multipathing
FCAL	Fibre Channel Arbitrated Loop	**IQN**	iSCSI Qualified Name
FDX	Full Duplex	**ISC**	Integrated Solutions Console
FFDC	First Failure Data Capture	**ISM**	IBM Support Management
FLOP	Floating Point Operation	**ISSO**	Information System Security Officer
FRU	Field Replaceable Unit		
FTP	File Transfer Protocol	**ISV**	Independent Software Vendor
GB	Gigabyte	**IT**	information technology
GID	Group ID	**ITM**	IBM Tivoli Monitoring
GPFS	General Parallel File System	**ITSO**	International Technical Support Organization
GSS	General Security Services		
GUI	Graphical User Interface	**ITUAM**	IBM Tivoli Usage and Accounting Manager
HACMP/PowerHA	High Availability Cluster Multiprocessing		
		IVM	Integrated Virtualization Manager
HBA	Host Bus Adapters		
HMC	Hardware Management Console	**JDBC**	Java Database Connectivity
		JFS	Journaled File System
HPC	High Performance Computing	**JFS2**	Journaled File System2
HPM	Hardware Performance Monitor	**KAT**	Kernel Authorization Table
		KB	Kilobytes
HTML	Hypertext Markup Language	**KCT**	Kernel Command Table
HTTP	Hypertext Transfer Protocol	**KDT**	Kernel Device Table
Hz	Hertz	**KRT**	Kernel Role Table
		KST	Kernel Security Table

LA	Link Aggregation	**NIM**	Network Installation Management
LACP	Link Aggregation Control Protocol	**NIMOL**	NIM on Linux
LAN	Local Area Network	**NIS**	Network Information Server
LDAP	Light Weight Directory Access Protocol	**NVRAM**	Non-Volatile Random Access Memory
LFS	Logical File System	**ODBC**	Open Database Connectivity
LFT	Low Function Terminal	**ODM**	Object Data Manager
LMB	Logical Memory Block	**OSGi**	Open Services Gateway Initiative
LOB	Line of Business		
LPAR	Logical Partition	**OSPF**	Open Shortest Path First
LPP	Licensed Program Product	**PCI**	Peripheral Component Interconnect
LPS	Limiting Privilege Set		
LUN	Logical Unit Number	**PID**	Process ID
LUNs	Logical Unit Numbers	**PIT**	Point-in-time
LV	Logical Volume	**PKI**	Public Key Infrastructure
LVCB	Logical Volume Control Block	**PLM**	Partition Load Manager
LVM	Logical Volume Manager	**PM**	Performance Monitor
LWI	Lightweight Infrastructure	**POSIX**	Portable Operating System Interface
MAC	Media Access Control	**POST**	Power-On Self-test
Mbps	Megabits Per Second	**POWER®**	Performance Optimization with Enhanced RISC (Architecture)
MBps	Megabytes Per Second		
MCM	Multichip Module		
ME	Management Edition	**PP**	Physical Partition
MIBs	Management Information Bases	**PPC**	Physical Processor Consumption
ML	Maintenance Level	**PTF**	Program Temporary Fix
MLS	Multi Level Security	**PTX**	Performance Toolbox
MP	Multiprocessor	**PV**	Physical Volume
MPIO	Multipath I/O	**PVID**	Physical Volume Identifier
MPS	Maximum Privilege Set	**QoS**	Quality of Service
MTU	Maximum Transmission Unit	**PAC**	Proactive Analysis Component
NDAF	Network Data Administration Facility	**RAID**	Redundant Array of Independent Disks
NFS	Network File System		
NIB	Network Interface Backup	**RAM**	Random Access Memory

RAS	Reliability, Availability, and Serviceability	**SMTP**	Simple Mail Transfer Protocol
RBAC	Role Based Access Control	**SNMP**	Simple Network Management Protocol
RCP	Remote Copy	**SO**	System Operator
RDAC	Redundant Disk Array Controller	**SP**	Service Processor
RIP	Routing Information Protocol	**SPOT**	Shared Product Object Tree
RISC	Reduced Instruction-Set Computer	**SQL**	Structured Query Language
RMC	Resource Monitoring and Control	**SRC**	System Resource Controller
		SRN	Service Request Number
ROI	Return On Investment	**SSA**	Serial Storage Architecture
RPC	Remote Procedure Call	**SSH**	Secure Shell
RPL	Remote Program Loader	**SSL**	Secure Sockets Layer
RPM	Red Hat Package Manager	**SUID**	Set User ID
RSA	Rivet, Shamir, Adelman	**SVC**	SAN Virtualization Controller
RSCT	Reliable Scalable Cluster Technology	**TADDM**	Tivoli
RSH	Remote Shell	**TCB**	Trusted Computing Base
RTE	Runtime Error	**TCP/IP**	Transmission Control Protocol/Internet Protocol
RUID	Real User ID	**TE**	Trusted Execution
S	System Scope	**TEM**	Tivoli Enterprise Monitor
SA	System Administrator	**TEMA**	Tivoli Enterprise Monitor Agent
SAN	Storage Area Network	**TEMS**	Tivoli Enterprise Monitoring Server
SCP	secure copy protocol		
SCSI	Small Computer System Interface	**TEP**	Trusted Execution Path
		TEPS	Tivoli Enterprise Portal Server
SDD	Subsystem Device Driver	**TERMS**	Tivoli Enterprise Remote Monitor Server
SED	Stack Execution Disable		
SLs	Sensitivity Labels	**TPM**	Tivoli Productivity Manager
SMI	Structure of Management Information	**TSA**	Tivoli System Automation
		TSD	Trusted Signature Database
SMIT	Systems Management Interface Tool	**TTL**	Time-to-live
		TUAM	Tivoli Usage and Accounting Manager
SMP	Symmetric Multiprocessor		
SMS	System Management Services	**UCS**	Universal-Coded Character Set
SMT	Simultaneous multi-threading	**UDB**	Universal Database

UID	User ID
VE	Virtualization Edition
VG	Volume Group
VIO	Virtual I/O
VIPA	Virtual IP Address
VLAN	Virtual Local Area Network
VMM	Virtual Memory Manager
VP	Virtual Processor
VPA	Visual Performance Analyzer
VPD	Vital Product Data
VPN	Virtual Private Network
VPSS	Variable Page Size Support
VRRP	Virtual Router Redundancy Protocol
VSD	Virtual Shared Disk
WAS	WebSphere Application Server
WED	WebSphere Everyplace® Deployment V6.0
WLM	Workload Manager
WPAR	Workload Partitions
WPS	Workload Partition Privilege Set

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 publications

For information about ordering these publications, see "How to get IBM Redbooks publications" on page 278. Note that several 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

Other publications

These publications are also relevant as further information sources:

► Documentation available on the support and services Web site includes:

  – User guides
  – System management guides
  – Application programmer guides
  – All commands reference volumes
  – Files reference
  – Technical reference volumes that are 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

How to get IBM Redbooks publications

You can search for, view, or download IBM Redbooks publications, Redpapers, Technotes, draft publications, and additional materials, as well as order hardcopy IBM Redbooks publications, 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

IBM AIX Enterprise Edition System Administration Guide

IBM AIX Enterprise Edition
System Administration Guide

Learn about AIX Enterprise Edition based on AIX Version 6.1

Obtain useful planning information

Explore an offering that combines AIX with IBM Tivoli Systems Management

This IBM Redbooks publication provides AIX Systems Administrators a broad understanding of the new features and products that make up the AIX Enterprise Edition offering. It will help you design or create a solution to maximize efficiencies and help organizations get the most from their Power Systems platform.

AIX Enterprise Edition includes all the features of AIX with additional management software that is designed to complement the features of AIX V6.1. AIX Enterprise Edition is designed to improve the manageability of a virtualized AIX environment. Potential benefits include:

► Improved service availability through access to relevant, real-time information about IT resources
► Enhanced operational efficiency though the visualization of the relationships of IT resources and applications
► Improved operational efficiency by centralizing monitoring and automating the response to service issues
► Better utilization of IT resources through consolidated views of virtualized resource usage
► Improved service availability through Live Application Mobility
► Simplified ordering of AIX and essential service management tools at an attractive price